

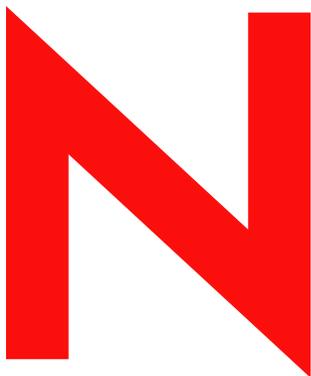
Novell ZENworks® Server Management

7

www.novell.com

ADMINISTRATION GUIDE

September 7, 2007



Novell®

Legal Notices

Novell, Inc. makes no representations or warranties with respect to the contents or use of this documentation, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to revise this publication and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc. makes no representations or warranties with respect to any software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to make changes to any and all parts of Novell software, at any time, without any obligation to notify any person or entity of such changes.

Any products or technical information provided under this Agreement may be subject to U.S. export controls and the trade laws of other countries. You agree to comply with all export control regulations and to obtain any required licenses or classification to export, re-export, or import deliverables. You agree not to export or re-export to entities on the current U.S. export exclusion lists or to any embargoed or terrorist countries as specified in the U.S. export laws. You agree to not use deliverables for prohibited nuclear, missile, or chemical biological weaponry end uses. Please refer to www.novell.com/info/exports/ for more information on exporting Novell software. Novell assumes no responsibility for your failure to obtain any necessary export approvals.

Copyright © 1999-2006 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

Novell, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at the Novell Legal Web site (<http://www.novell.com/company/legal/patents>) and one or more additional patents or pending patent applications in the U.S. and in other countries.

Novell, Inc.
404 Wyman Street, Suite 500
Waltham, MA 02451
U.S.A.
www.novell.com

Online Documentation: To access the online documentation for this and other Novell products, and to get updates, see www.novell.com/documentation.

Novell Trademarks

ConsoleOne is a registered trademark of Novell, Inc., in the United States and other countries.

eDirectory is a trademark of Novell, Inc.

GroupWise is a registered trademark of Novell, Inc., in the United States and other countries.

IPX is a trademark of Novell, Inc.

NCP is a trademark of Novell, Inc.

NDS is a registered trademark of Novell, Inc., in the United States and other countries.

NetExplorer is a trademark of Novell, Inc.

NetWare is a registered trademark of Novell, Inc., in the United States and other countries.

NetWare Core Protocol is a trademark of Novell, Inc.

NetWare Management Agent is a trademark of Novell, Inc.

NetWare SFT III is a trademark of Novell, Inc.

NLM is a trademark of Novell, Inc.

Novell is a registered trademark of Novell, Inc., in the United States and other countries.

Novell Client is a trademark of Novell, Inc.

Novell Technical Services is a service mark of Novell, Inc.

SPX is a trademark of Novell, Inc.

SUSE is a registered trademark of Novell, Inc., in the United States and other countries.

ZENworks is a registered trademark of Novell, Inc., in the United States and other countries.

Third-Party Materials

All third-party trademarks are the property of their respective owners.

Contents

| | |
|--|-----------|
| About This Guide | 25 |
| Part I Policy and Distribution Services | 27 |
| 1 Post-Installation Setup | 29 |
| 1.1 Planning Your Distribution System | 29 |
| 1.1.1 Overview | 30 |
| 1.1.2 Selecting Your Distributions | 32 |
| 1.1.3 Understanding Your Network Topology | 36 |
| 1.1.4 Are Additional Distributors Needed? | 37 |
| 1.1.5 Other Subscribers To Be Installed? | 41 |
| 1.1.6 Determining the Distribution Flow | 41 |
| 1.1.7 Understanding Distribution Security | 44 |
| 1.1.8 Determining the Channels for the Distributions | 46 |
| 1.1.9 Determining Subscribers' Subscriptions | 47 |
| 1.1.10 Determining the Distribution Schedules | 48 |
| 1.2 Configuring Your Distribution System | 50 |
| 1.2.1 Installing Additional Distributors, Databases, and Subscribers | 50 |
| 1.2.2 Setting Up Additional Distribution Security | 54 |
| 1.2.3 Configuring the Distribution Flow | 55 |
| 1.2.4 Creating the Distributions and Related Channels | 57 |
| 1.2.5 Subscribing to the Distributions | 59 |
| 1.2.6 Sending the Distributions | 60 |
| 1.3 Managing Your Distribution System | 60 |
| 2 Novell iManager | 63 |
| 2.1 Accessing the ZENworks Server Management Role in iManager | 63 |
| 2.1.1 Logging in to iManager | 64 |
| 2.1.2 Becoming Familiar with the Interface | 65 |
| 2.1.3 Viewing the Roles and Tasks | 65 |
| 2.2 Managing Tiered Electronic Distribution Objects | 67 |
| 2.2.1 Creating Tiered Electronic Distribution Objects in iManager | 68 |
| 2.2.2 Editing Tiered Electronic Distribution Object Properties in iManager | 68 |
| 2.2.3 Deleting Tiered Electronic Distribution Objects in iManager | 69 |
| 2.3 Monitoring the Distribution Process | 69 |
| 2.3.1 Using the Tiered Distribution View | 69 |
| 2.3.2 Using the Subscriber Distribution View | 70 |
| 2.4 Managing the Agents through Remote Web Console | 72 |
| 2.4.1 Setting Up Passwords for Remote Web Console | 73 |
| 2.4.2 Managing the Distributor Agent | 77 |
| 2.4.3 Managing the Policy/Package Agent | 80 |
| 2.4.4 Opening Multiple Remote Web Console Windows | 81 |
| 2.5 Comparing the ZENworks Server Management Role in iManager with ConsoleOne Capabilities | 82 |
| 3 Tiered Electronic Distribution | 85 |
| 3.1 Common Distribution Tasks | 85 |

| | | |
|--------|---|-----|
| 3.2 | Understanding Tiered Electronic Distribution | 87 |
| 3.2.1 | Distribution Management through Tiered Electronic Distribution | 88 |
| 3.2.2 | The Basic Distribution Process | 88 |
| 3.2.3 | Tiered Electronic Distribution's eDirectory Objects | 89 |
| 3.2.4 | Relationships of the Tiered Electronic Distribution Objects | 89 |
| 3.2.5 | Physical Network Connections | 90 |
| 3.2.6 | Distribution Flow Details | 90 |
| 3.2.7 | Tiered Electronic Distribution Processes | 91 |
| 3.2.8 | The Tiered Distribution Model | 93 |
| 3.2.9 | Tiered Electronic Distribution's Key Components | 94 |
| 3.3 | Distributors | 95 |
| 3.3.1 | Understanding Distributors | 95 |
| 3.3.2 | Understanding Distribution Routing | 97 |
| 3.3.3 | Creating Distributors | 106 |
| 3.3.4 | Configuring Distributors | 106 |
| 3.3.5 | Manually Refreshing the Distributor | 109 |
| 3.3.6 | Deleting a Distributor Object and How Its Distributions Are Affected | 110 |
| 3.4 | Distributions | 110 |
| 3.4.1 | Understanding Distributions | 110 |
| 3.4.2 | Distribution Issues | 114 |
| 3.4.3 | Determining the Distributions | 117 |
| 3.4.4 | Creating a Distribution | 123 |
| 3.4.5 | Prioritizing Distributions | 126 |
| 3.4.6 | Pre and Post Processing for Distributions | 126 |
| 3.4.7 | Reassigning a Distribution to Another Distributor | 133 |
| 3.4.8 | Deleting a Distribution | 136 |
| 3.4.9 | Removing a Distribution Object - Auto Removal of Temporary Files | 137 |
| 3.4.10 | Handling Orphaned Distributions | 139 |
| 3.4.11 | Manually Importing and Exporting Distributions | 141 |
| 3.4.12 | Using the Distribution Wizard | 143 |
| 3.5 | Channels | 144 |
| 3.5.1 | Understanding Channels | 145 |
| 3.5.2 | Creating and Configuring Channels | 146 |
| 3.5.3 | Forcing a Channel To Be Sent | 147 |
| 3.6 | Subscribers | 147 |
| 3.6.1 | Understanding Subscribers | 148 |
| 3.6.2 | Creating Subscribers | 149 |
| 3.6.3 | Configuring Subscribers | 150 |
| 3.6.4 | Updating Subscriber Configurations | 153 |
| 3.6.5 | Associating Subscribers with Channels | 154 |
| 3.6.6 | Deleting Subscriber Objects That Are Part of a Distributor's Routing Hierarchy | 155 |
| 3.7 | Subscriber Groups | 155 |
| 3.7.1 | Understanding Subscriber Groups | 155 |
| 3.7.2 | Creating and Configuring Subscriber Groups | 156 |
| 3.8 | External Subscribers | 156 |
| 3.8.1 | Understanding External Subscribers | 156 |
| 3.8.2 | Using External Subscribers for Out-of-Tree Distributions | 162 |
| 3.8.3 | Creating and Configuring External Subscribers | 164 |
| 3.9 | Configuring Multiple Tiered Electronic Distribution Objects | 165 |
| 3.9.1 | Issues with Modifying Multiple Tiered Electronic Distribution Object Properties | 166 |
| 3.9.2 | Modifying Multiple Tiered Electronic Distribution Object Properties | 167 |
| 3.9.3 | Property Tabs Available for Multiple-Object Modifications | 167 |
| 3.10 | Sending Distributions | 172 |
| 3.10.1 | Understanding the Distribution Processes | 172 |
| 3.10.2 | Forcing a Single Distribution To Be Sent | 173 |
| 3.10.3 | Sending Distributions Through Parent Subscribers | 174 |
| 3.10.4 | Sending Distributions between Trees | 174 |

| | | |
|----------|---|------------|
| 3.10.5 | Sending Distributions: Firewall and Cluster Issues | 175 |
| 3.11 | Miscellaneous Tiered Electronic Distribution Issues | 175 |
| 3.11.1 | Directory Sync Granularity for File Distributions | 176 |
| 3.11.2 | Understanding Dependencies in Tiered Electronic Distribution | 181 |
| 3.11.3 | System Resources and Server Behavior | 181 |
| 3.11.4 | Controlling I/O Rates and Concurrent Distributions. | 182 |
| 3.11.5 | Minimizing Messaging Traffic | 183 |
| 3.11.6 | Changing DNS Names or IP Addresses for Tiered Electronic Distribution Servers | 185 |
| 3.11.7 | When a Tiered Electronic Distribution Process Fails. | 185 |
| 3.12 | Working Directories | 186 |
| 3.12.1 | NetWare Distributor Directories | 186 |
| 3.12.2 | NetWare Subscriber Directories | 188 |
| 3.12.3 | Windows Distributor Directories | 189 |
| 3.12.4 | Windows Subscriber Directories | 189 |
| 3.12.5 | Linux or Solaris Distributor Directories | 190 |
| 3.12.6 | Linux or Solaris Subscriber Directories | 190 |
| 3.13 | Editing the Tednode.properties File. | 190 |
| 4 | Server Policies | 193 |
| 4.1 | Understanding Server Policies | 193 |
| 4.1.1 | Configuration and Behavioral Management through Server Policies | 193 |
| 4.1.2 | Server Policies and Policy Packages | 194 |
| 4.1.3 | Policy Characteristics | 195 |
| 4.1.4 | Server Policies Architecture | 196 |
| 4.1.5 | Enforcing Policies | 199 |
| 4.1.6 | Server Policy Descriptions | 200 |
| 4.2 | Creating a Policy Package. | 204 |
| 4.2.1 | Creating a Policies Container | 205 |
| 4.2.2 | Creating a Policy Package Object. | 205 |
| 4.3 | Configuring Server Policies | 205 |
| 4.3.1 | Compiling Zentrap.mib | 206 |
| 4.3.2 | Configuring the Container Package Policy | 206 |
| 4.3.3 | Configuring Service Location Package Policies | 207 |
| 4.3.4 | Configuring Distributed Server Package Policies | 214 |
| 4.3.5 | Creating Custom Log Files Using Policies | 232 |
| 4.4 | Enabling Policies | 233 |
| 4.5 | Distributing Policies | 233 |
| 4.6 | Associating Policies | 233 |
| 4.6.1 | Associating a Policy Package to the Distributor Object. | 234 |
| 4.6.2 | Associating the Distributor Object to a Policy Package. | 234 |
| 4.7 | Scheduling Policies | 234 |
| 4.8 | Viewing Effective Policies | 234 |
| 4.9 | Changing Policy Enforcement | 234 |
| 4.9.1 | Modifying a Policy That Is Being Enforced | 235 |
| 4.9.2 | Stopping a Specific Policy From Being Enforced | 235 |
| 4.9.3 | Removing Policy Enforcement for a Specific Subscriber | 235 |
| 4.9.4 | Stopping Enforcement of a Policy Package Distribution | 236 |
| 5 | Server Software Packages | 237 |
| 5.1 | Software Management through Server Software Packages | 237 |
| 5.2 | Understanding Server Software Packages | 237 |
| 5.2.1 | Understanding Server Software Packages and Components | 238 |
| 5.2.2 | Understanding Software Package and Component Configurations | 238 |
| 5.2.3 | Determining the Installation Order of Software Packages. | 239 |

| | | |
|----------|---|------------|
| 5.2.4 | Executing Extracted Files | 240 |
| 5.2.5 | Compiling Software Packages | 241 |
| 5.2.6 | Accessing Software Packages | 241 |
| 5.2.7 | Distributing Software Packages | 242 |
| 5.2.8 | Distributing Software Packages to a Cluster | 242 |
| 5.2.9 | Managing Server Software Packages | 243 |
| 5.2.10 | Failure of Software Package Installations | 248 |
| 5.2.11 | Rolling Back Software Package Installations | 249 |
| 5.3 | Planning Server Software Packages | 249 |
| 5.3.1 | Which Files or Applications Do I Want to Distribute? | 250 |
| 5.3.2 | What Software Package Components Are Needed? | 250 |
| 5.3.3 | What Minimum Requirements Are Needed? | 250 |
| 5.4 | Setting Up Server Software Packages | 251 |
| 5.4.1 | Setting Up Multiple-Workstation Management for Server Software Packages | 251 |
| 5.4.2 | Creating a Server Software Package | 255 |
| 5.4.3 | Configuring the Server Software Package | 255 |
| 5.4.4 | Creating the Software Package Components | 256 |
| 5.4.5 | Configuring the Software Package Components | 256 |
| 5.4.6 | Compiling a Software Package | 267 |
| 5.4.7 | Distributing the Software Package | 267 |
| 5.5 | Using Server Software Packages to Delete Directories on Servers | 268 |
| 5.5.1 | Setting Up Variables for Use With the Server Software Package | 268 |
| 5.5.2 | Creating the Server Software Package | 269 |
| 5.5.3 | Creating and Configuring the Server Software Package Component | 269 |
| 5.5.4 | Compiling the Server Software Package | 271 |
| 5.5.5 | Manually Testing that the Directories Have Been Deleted | 271 |
| 5.5.6 | What's Next | 271 |
| 6 | Desktop Application Distribution | 273 |
| 6.1 | Understanding Desktop Application Distributions | 273 |
| 6.1.1 | The Purpose of Desktop Application Distributions | 273 |
| 6.1.2 | Distributed Application Issues | 275 |
| 6.1.3 | Miscellaneous Issues | 286 |
| 6.2 | Requirements | 288 |
| 6.3 | Creating a Desktop Application Distribution | 289 |
| 6.3.1 | Understanding the Desktop Application Distribution Wizard | 289 |
| 6.3.2 | Creating the Distribution | 290 |
| 6.4 | Rebuilding Desktop Application Distributions | 296 |
| 6.4.1 | All Attributes Are Updated | 296 |
| 6.4.2 | Triggering a Rebuild | 296 |
| 6.5 | Cleaning Up Desktop Application Distribution Files | 297 |
| 6.6 | Sending Desktop Application Distributions Tree-To-Tree | 298 |
| 7 | Security in Policy and Distribution Services | 299 |
| 7.1 | Distribution Security Using Signed Certificates and Digests | 299 |
| 7.1.1 | Understanding Digests | 300 |
| 7.1.2 | Understanding Certificate Usage in Policy and Distribution Services | 300 |
| 7.1.3 | Important Points about Certificates | 301 |
| 7.1.4 | ConsoleOne User Rights and Certificate Copying | 302 |
| 7.1.5 | Certificate File Locations | 303 |
| 7.1.6 | Resolving Certificates | 303 |
| 7.1.7 | Handling Invalid Certificates | 304 |
| 7.1.8 | Certificate and Private Key Directories | 307 |
| 7.1.9 | Creating Security Certificates for Non-Encrypted Distributions | 307 |
| 7.1.10 | Manually Copying Certificates for Non-Encrypted Distributions | 308 |

| | | |
|-----------|---|------------|
| 7.2 | Distribution Security Using Encryption | 309 |
| 7.2.1 | Creating and Copying Encryption Certificates | 309 |
| 7.2.2 | Sending an Encrypted Distribution | 311 |
| 7.2.3 | Extracting an Encrypted Distribution | 312 |
| 7.3 | Security for Inter-Server Communication Across Non-Secured Connections | 313 |
| 7.3.1 | Terms Used in This Section | 313 |
| 7.3.2 | Security Certificates | 314 |
| 7.3.3 | Using SSL | 314 |
| 7.3.4 | Format of the Password File | 314 |
| 7.3.5 | TCP/IP Addresses and DNS Names | 315 |
| 8 | Scheduling | 317 |
| 8.1 | Understanding Scheduling in Policy and Distribution Services | 317 |
| 8.1.1 | Why Scheduling is Necessary for Distributions | 317 |
| 8.1.2 | Scheduling Is Required for Some Server Policies | 318 |
| 8.1.3 | Scheduling Differences Between Server Policies and Tiered Electronic Distribution | 318 |
| 8.1.4 | Precedence of the Tiered Electronic Distribution Policy | 319 |
| 8.2 | Scheduling and Tiered Electronic Distribution Objects | 319 |
| 8.2.1 | Understanding the Tiered Electronic Distribution Objects and Their Schedules | 319 |
| 8.2.2 | How the Tiered Electronic Distribution Schedules Interrelate | 324 |
| 8.2.3 | The Three Timing Aspects of Scheduling | 326 |
| 8.2.4 | Approaches to Scheduling | 329 |
| 8.2.5 | Scheduling Issues | 331 |
| 8.3 | Scheduling and Server Policies | 338 |
| 8.3.1 | Policy Schedules Versus Distribution Schedules | 338 |
| 8.3.2 | Scheduling a Server Policy | 339 |
| 8.3.3 | Editing the Default Package Schedule | 339 |
| 9 | Variables | 341 |
| 9.1 | Understanding Variables | 341 |
| 9.1.1 | Why Variables? | 341 |
| 9.1.2 | Variable Usage | 342 |
| 9.1.3 | Variable Usage Differences | 343 |
| 9.1.4 | Precedence for Determining Which Variable to Use | 344 |
| 9.1.5 | Distribution Variable Example | 344 |
| 9.2 | Types of Variables | 344 |
| 9.2.1 | Predefined Variables | 344 |
| 9.2.2 | User-Defined Variables | 346 |
| 9.3 | Defining a Variable | 346 |
| 9.3.1 | Defining Default Variables for All Subscribers | 347 |
| 9.3.2 | Defining Variables for a Specific Subscriber | 347 |
| 9.3.3 | Defining Variables for a Server Software Package | 348 |
| 9.4 | Viewing All Variables in iManager | 348 |
| 9.5 | Using a Variable to Change a Subscriber's Console Prompt | 348 |
| 9.6 | Using Variables to Control File Extraction | 349 |
| 10 | ZENworks Database | 351 |
| 10.1 | Understanding the ZENworks Database | 351 |
| 10.1.1 | The Database Engine | 351 |
| 10.1.2 | The Database File | 351 |
| 10.1.3 | The Database Object | 351 |
| 10.1.4 | Running the Database | 352 |

| | | |
|--------|---|-----|
| 10.1.5 | Database Caching | 352 |
| 10.1.6 | Database Information | 352 |
| 10.1.7 | Coexisting Databases | 353 |
| 10.2 | Determining How Many Databases You Need | 353 |
| 10.2.1 | Database Logging and Tiered Electronic Distribution Reporting | 354 |
| 10.2.2 | Multiple Databases | 355 |
| 10.3 | Installing and Connecting to the Server Management Database | 356 |
| 10.3.1 | Installing the Database | 357 |
| 10.3.2 | Connecting to the Database | 358 |
| 10.4 | Creating a ZENworks Database Object | 359 |
| 10.5 | Purging the Database | 359 |
| 10.5.1 | Tiered Electronic Distribution Information | 360 |
| 10.5.2 | Server Policies Information | 360 |

11 Reporting 363

| | | |
|--------|--|-----|
| 11.1 | Understanding Policy and Distribution Services Reporting | 363 |
| 11.1.1 | Reporting Categories | 363 |
| 11.1.2 | Reporting Scope | 364 |
| 11.1.3 | Accessing Reports | 364 |
| 11.1.4 | Creating and Storing Report Information | 364 |
| 11.2 | Report Descriptions | 365 |
| 11.2.1 | Tiered Electronic Distribution Reports | 365 |
| 11.2.2 | Server Policy Reports | 367 |
| 11.3 | Generating Reports | 369 |
| 11.4 | Creating Customized Reports | 370 |
| 11.4.1 | Default Sybase Database User ID and Password | 370 |
| 11.4.2 | Server Policies Database Contents | 371 |
| 11.4.3 | Tiered Electronic Distribution Database Contents | 378 |

A Distribution Types 383

| | | |
|-------|-----------------------------|-----|
| A.1 | Desktop Application | 383 |
| A.2 | File | 383 |
| A.2.1 | Files to Be Distributed | 384 |
| A.2.2 | New Target | 384 |
| A.2.3 | Add Directory | 385 |
| A.2.4 | Add Files | 385 |
| A.2.5 | Delete | 385 |
| A.2.6 | Synchronize/Desynchronize | 386 |
| A.2.7 | Verify Distributions | 386 |
| A.2.8 | Maintain Trustees | 387 |
| A.2.9 | Extract Error Handling | 387 |
| A.3 | FTP | 387 |
| A.3.1 | Files To Be Distributed | 388 |
| A.3.2 | New FTP Source | 388 |
| A.3.3 | New Target | 388 |
| A.3.4 | Add Directory | 389 |
| A.3.5 | Add Files | 389 |
| A.3.6 | Delete | 389 |
| A.3.7 | Properties | 390 |
| A.3.8 | Binary Transfer | 390 |
| A.3.9 | Include Symbolic Link Files | 390 |
| A.4 | HTTP | 390 |
| A.4.1 | Files To Be Distributed | 390 |
| A.4.2 | New Target | 391 |
| A.4.3 | Add Directory | 391 |

| | | |
|--------------------------------------|---|----------------|
| A.4.4 | Add Files | 391 |
| A.4.5 | Delete | 391 |
| A.5 | MSI | 392 |
| A.5.1 | Adding | 392 |
| A.5.2 | Removing | 392 |
| A.5.3 | Configuring | 392 |
| A.5.4 | Rearranging | 395 |
| A.6 | Policy Package | 395 |
| A.6.1 | Up/Down | 395 |
| A.6.2 | Add | 395 |
| A.6.3 | Delete | 395 |
| A.6.4 | Properties | 395 |
| A.6.5 | The Following Policy Packages Will Be Distributed | 395 |
| A.7 | RPM | 396 |
| A.7.1 | Up/Down | 396 |
| A.7.2 | Add From Distributor | 396 |
| A.7.3 | Add From FTP Site | 396 |
| A.7.4 | Delete | 396 |
| A.7.5 | Selected Packages | 396 |
| A.7.6 | Installation Parameters | 396 |
| A.8 | Software Package | 397 |
| A.8.1 | Up/Down | 397 |
| A.8.2 | Add | 397 |
| A.8.3 | Delete | 397 |
| A.8.4 | Selected Software Packages | 397 |
| B Schedule Types | | 399 |
| B.1 | Daily | 400 |
| B.2 | Event | 400 |
| B.3 | Interval | 400 |
| B.4 | Monthly | 401 |
| B.5 | Never | 401 |
| B.6 | Package Schedule | 401 |
| B.7 | Relative | 402 |
| B.8 | Run Immediately | 402 |
| B.9 | Time | 402 |
| B.10 | Weekly | 403 |
| B.11 | Yearly | 403 |
| C Server Console Commands | | 405 |
| C.1 | ZENworks Server Management Console Commands | 405 |
| C.2 | Java Console Commands | 408 |
| D Load/Unload Actions | | 411 |
| D.1 | Load NLM/Process | 411 |
| D.2 | Load Java Class | 411 |
| D.3 | Unload Process | 412 |
| D.4 | Start Service | 412 |
| D.5 | Stop Service | 412 |

| | | |
|----------------|--|------------|
| E | Requirements for Server Software Packages | 413 |
| E.1 | Operating System | 413 |
| E.2 | Memory (RAM) | 416 |
| E.3 | Disk Space | 416 |
| E.4 | SET Commands | 417 |
| E.5 | Registry | 418 |
| E.6 | File | 418 |
| E.7 | Products.dat | 418 |
| F | Registry Entries for Server Software Package Components | 421 |
| F.1 | Key | 421 |
| F.2 | Binary | 422 |
| F.3 | Expand String | 422 |
| F.4 | (Default) | 422 |
| F.5 | DWord | 423 |
| F.6 | Multi-Value String | 423 |
| F.7 | String | 423 |
| G | Client Access in Linux | 425 |
| G.1 | Using Samba | 425 |
| G.2 | Using NCP Shares | 425 |
| H | Configuration Planning Worksheet | 427 |
| I | Documentation Updates | 437 |
| I.1 | March 29, 2007 | 437 |
| I.2 | August 16, 2006 | 438 |
| I.2.1 | Load/Unload Actions | 438 |
| I.2.2 | Security in Policy and Distribution Services | 438 |
| I.2.3 | Server Policies | 439 |
| I.2.4 | Tiered Electronic Distribution | 439 |
| I.3 | July 14, 2006 (Support Pack 1) | 439 |
| I.3.1 | Client Access in Linux | 439 |
| I.3.2 | Desktop Application Distribution | 439 |
| I.3.3 | Distribution Types | 440 |
| I.3.4 | Novell iManager | 440 |
| I.3.5 | Requirements for Server Software Packages | 440 |
| I.3.6 | Schedule Types | 441 |
| I.3.7 | Server Policies | 441 |
| I.3.8 | Tiered Electronic Distribution | 441 |
| I.4 | December 9, 2005 | 441 |
| I.5 | October 7, 2005 | 442 |
| I.5.1 | Desktop Application Distribution | 442 |
| Part II | Server Inventory | 443 |
| 12 | Understanding Server Inventory | 445 |
| 12.1 | Server Inventory Terminology | 445 |
| 12.2 | Overview of Server Inventory Components | 446 |

| | | |
|-----------|--|------------|
| 12.2.1 | Inventory Scanners | 446 |
| 12.2.2 | Inventory Components on Inventory Servers | 447 |
| 12.2.3 | Inventory Database | 448 |
| 12.2.4 | Management Console | 448 |
| 12.3 | Understanding Inventory Scanning Cycle | 448 |
| 12.4 | Understanding the Inventory Server Roles | 448 |
| 12.4.1 | Root Server | 449 |
| 12.4.2 | Root Server with Inventoried Servers | 450 |
| 12.4.3 | Intermediate Server | 450 |
| 12.4.4 | Intermediate Server with Database | 451 |
| 12.4.5 | Intermediate Server with Inventoried Servers | 452 |
| 12.4.6 | Intermediate Server with Database and Inventoried Servers | 453 |
| 12.4.7 | Leaf Server | 454 |
| 12.4.8 | Leaf Server with Database | 455 |
| 12.4.9 | Standalone Server | 456 |
| 12.4.10 | Quick Reference Table of the Inventory Server Roles | 457 |
| 13 | Setting Up Server Inventory | 459 |
| 13.1 | Deploying Server Inventory | 459 |
| 13.1.1 | Simple Deployment | 459 |
| 13.1.2 | Advanced Deployment | 463 |
| 13.1.3 | Understanding the Effects of Server Inventory Installation | 480 |
| 13.1.4 | Starting and Stopping the Inventory Service | 482 |
| 13.1.5 | Changing the Role of the Inventory Server | 483 |
| 13.2 | Setting Up the Inventory Database | 493 |
| 13.2.1 | Setting Up the Sybase Inventory Database | 493 |
| 13.2.2 | Setting Up the Oracle Inventory Database | 500 |
| 13.2.3 | Setting Up the MS SQL Server 2000 or MS SQL Server 2005 Inventory Database | 510 |
| 13.3 | Configuring the Inventory Service Object | 520 |
| 13.4 | Configuring the Database Location Policy | 521 |
| 13.5 | Configuring the Server Inventory Policy | 522 |
| 13.6 | Configuring the Roll-Up Policy | 524 |
| 13.7 | Configuring the Dictionary Update Policy | 525 |
| 13.8 | Setting Up Distribution of Dictionary | 526 |
| 14 | Understanding the Server Inventory Components | 529 |
| 14.1 | Understanding the Inventory Service Manager | 529 |
| 14.1.1 | List of Services | 529 |
| 14.1.2 | Services on NetWare Inventory Servers | 530 |
| 14.1.3 | Services on Windows Inventory Servers | 530 |
| 14.2 | Understanding the Server Configuration Service | 532 |
| 14.3 | Understanding the Inventory Scanner | 532 |
| 14.3.1 | Inventory Scanning Process | 532 |
| 14.3.2 | Scanning for the NetWare Inventoried Servers | 533 |
| 14.3.3 | Scanning for the Windows Inventoried Servers | 535 |
| 14.4 | Understanding the Sender and Receiver | 539 |
| 14.4.1 | Understanding the Sender | 541 |
| 14.4.2 | Understanding the Receiver | 542 |
| 14.4.3 | Understanding the Compressed Scan Data File | 542 |
| 14.4.4 | Sender-Receiver Directories | 543 |
| 14.5 | Understanding the Selector | 544 |
| 14.6 | Understanding the Storer | 545 |
| 14.7 | Understanding the Dictionary Provider and Dictionary Consumer | 546 |
| 14.8 | Understanding the Upgrade Service | 546 |

| | | |
|-------|---|-----|
| 14.9 | An Overview of the Inventory Components on the Inventory Server | 546 |
| 14.10 | Understanding the Inventory Database | 547 |

15 Understanding the ZENworks 7 Server Managements Inventory Database Schema 549

| | | |
|--------|---|-----|
| 15.1 | Overview. | 549 |
| 15.2 | CIM Schema | 550 |
| 15.2.1 | CIM-to-Relational Mapping | 553 |
| 15.3 | Inventory Database Schema in ZENworks 7 Server Management | 556 |
| 15.3.1 | Case Study of CIM Schema Implementation in ZENworks 7 Server Management | 557 |
| 15.3.2 | Legends for Schema Diagrams | 559 |
| 15.3.3 | Schema Diagrams of CIM and the Extension Schema in ZENworks 7 Server Management | 559 |
| 15.3.4 | Software Inventory Schema | 567 |
| 15.3.5 | Sample Inventory Database Queries | 573 |

16 Managing Your Inventory Information 581

| | | |
|---------|---|-----|
| 16.1 | Viewing the Inventory Servers Deployed for Inventory. | 581 |
| 16.2 | Customizing the Hardware Inventory Information To Be Scanned. | 582 |
| 16.2.1 | Scanning for Vendor-Specific Asset Information from DMI | 582 |
| 16.2.2 | Customizing the Hardware Scanning Information of Jaz and Zip Drive Vendors. | 583 |
| 16.2.3 | Customizing the Hardware Information for Monitor Size | 584 |
| 16.3 | Customizing the Software Inventory Information To Be Scanned For the ZENworks 7 Inventoried Servers | 586 |
| 16.3.1 | What is the ZENworks Software Dictionary? | 587 |
| 16.3.2 | What is a Software Dictionary Rule? | 588 |
| 16.3.3 | What is a Software Identifier? | 588 |
| 16.3.4 | What is a Key Identifier? | 588 |
| 16.3.5 | What is an Unidentified Software? | 588 |
| 16.3.6 | What is an Inherited Rule? | 588 |
| 16.3.7 | What is An Overriding Rule? | 588 |
| 16.3.8 | Understanding the Usage and Precedence of ZENworks Software Dictionary Rules | 589 |
| 16.3.9 | Understanding the Software Dictionary Pattern Types | 595 |
| 16.3.10 | Configuring the Software Dictionary Rules | 596 |
| 16.3.11 | Ignore Default File-Software Mapping Rules | 599 |
| 16.3.12 | Software Dictionary | 599 |
| 16.3.13 | Report Files with These File Extensions As Unidentified Software | 603 |
| 16.3.14 | Manage Unidentified Software | 604 |
| 16.3.15 | Report Multiple Software Versions | 605 |
| 16.3.16 | Report Disk Space Used by File Extensions | 607 |
| 16.3.17 | Software Scanning Filters - Drives and Directories | 608 |
| 16.3.18 | Software Scanning Filters - File Extensions. | 613 |
| 16.3.19 | Software Scanning Filters - Files | 615 |
| 16.3.20 | Software Scanning Filters - Software. | 616 |
| 16.3.21 | Disk Usage Scanning Filters - Drives and Directories | 618 |
| 16.3.22 | Disk Usage Scanning Filters - Files | 622 |
| 16.3.23 | Vendor Name Aliases | 623 |
| 16.3.24 | Software Name Aliases | 625 |
| 16.3.25 | Reconcile Software | 626 |
| 16.3.26 | Sorting Entries in the Table | 627 |
| 16.3.27 | Filtering Entries in the Table | 627 |
| 16.3.28 | Refreshing Entries in the Table | 628 |
| 16.3.29 | Disabling File Scan | 628 |
| 16.3.30 | Base-lining the Software Dictionary Deployment | 629 |

| | | |
|---|--|------------|
| 16.3.31 | Viewing Software Information in the Inventory Summary | 630 |
| 16.3.32 | Generating Software Inventory Reports | 630 |
| 16.4 | Customizing the Software Inventory Information To Be Scanned For ZENworks for Servers 3.x Inventoried Servers | 632 |
| 16.5 | Removing Redundant Inventoried Servers from the Inventory Database | 632 |
| 17 Viewing Inventory Information | | 635 |
| 17.1 | Viewing the Inventory Information Using ConsoleOne | 635 |
| 17.1.1 | Configuring the Inventory Database | 635 |
| 17.1.2 | Viewing the Inventory Summary of an Inventoried Server | 636 |
| 17.1.3 | Viewing Inventory Information of Inventoried Servers by Querying the Database | 649 |
| 17.1.4 | Running Inventory Reports | 651 |
| 17.1.5 | Quickly and Easily Viewing the Inventory Data Using Quick Reports | 664 |
| 17.2 | Exporting the Inventory Information | 676 |
| 17.2.1 | Procedure to Export the Inventory Information | 676 |
| 17.2.2 | Loading an Existing Configuration File | 679 |
| 17.2.3 | Running the Data Export Program from the Inventory Server | 680 |
| 17.2.4 | An Overview of XML and the Contents of an XML File | 681 |
| 17.3 | Retrieving Inventory information from the Inventory Database Without Using the CIM Schema | 683 |
| 17.3.1 | List of Inventory Views | 683 |
| 17.3.2 | How to Use the Inventory Views | 709 |
| 18 Monitoring Server Inventory Using Status Logs | | 711 |
| 18.1 | Viewing the Scan Status of an Inventoried Server | 711 |
| 18.2 | Viewing the Roll-Up History of the Inventory Server | 712 |
| 18.3 | Viewing the Status of Inventory Components on an Inventory Server | 712 |
| 18.4 | Viewing the Status of the Last Scan on the Inventoried Server | 713 |
| 18.5 | Viewing the Roll-Up Log for the Inventory Servers | 713 |
| 18.6 | Exporting the Inventory Status Log Files | 714 |
| 18.7 | Overview of Status Logs and Scan Logs | 714 |
| 18.8 | Viewing the Status Log in XML Format | 715 |
| J Performance Tips | | 717 |
| J.1 | Database Parameter Tuning Tips | 717 |
| J.1.1 | Sybase in the NetWare and Windows Environments | 717 |
| J.1.2 | Oracle in the NetWare, Windows, and Linux Environments | 719 |
| J.1.3 | Optimizing the Performance of the Oracle Database | 720 |
| J.1.4 | MS SQL in the Windows Environment | 721 |
| J.2 | Improving the Throughput of the Inventory Storer | 721 |
| J.2.1 | Factors to be Considered Before Deployment | 721 |
| J.2.2 | Procedure to Improve the Throughput | 722 |
| J.2.3 | Recommendations for Administering the ZENworks Inventory Server | 724 |
| J.2.4 | Recommendations for Administering the Inventory Database | 724 |
| J.3 | Performance Tips for the Inventory Server (Support Pack 1) | 724 |
| J.4 | Performance Tips for the Inventory ConsoleOne Utilities | 726 |
| J.4.1 | Inventory Reports Performance Tips | 727 |
| J.4.2 | Inventory Data Export Performance Tips | 727 |
| J.4.3 | Inventory Query Performance Tips | 727 |
| J.5 | References | 727 |

| | | |
|----------|---|------------|
| K | Hardware Information Collected by the Inventory Scanners | 729 |
| K.1 | Hardware Information Collected on NetWare Inventoried Servers | 729 |
| K.2 | Hardware Information Collected on Windows Inventoried Servers | 734 |
| | | |
| L | ZENworks 7 Server Management Inventory Attributes | 747 |
| | | |
| M | Enumeration Values | 767 |
| M.1 | Enumeration Values for General-System Information-Management Technology | 767 |
| M.2 | Enumeration Values for General-Inventory Information-Scan Mode | 768 |
| M.3 | Enumeration Values for Software-Operating Systems-Windows - Name. | 768 |
| M.4 | Enumeration Values for Installation Repository | 768 |
| M.5 | Enumeration Values for Hardware-Display Adapter-Video Architecture | 769 |
| M.6 | Enumeration Values for Hardware-Display Adapter-Video Memory Type | 769 |
| M.7 | Enumeration Values for Hardware-Pointing Device-Name. | 769 |
| M.8 | Enumeration Values for Hardware-Battery-Chemistry | 769 |
| M.9 | Enumeration Values for Hardware-Processor-Processor Family | 770 |
| M.10 | Enumeration Values for Hardware-Processor-Upgrade Method | 770 |
| M.11 | Enumeration Values for Hardware-Chassis-Chassis Type. | 770 |
| M.12 | Enumeration Values for Hardware-Bus-Protocol Supported | 771 |
| M.13 | Enumeration Values for Hardware-Processor-Role | 771 |
| M.14 | Enumeration Values for System-System Cache-Level. | 771 |
| M.15 | Enumeration Values for System-System Cache-Cache Type | 771 |
| M.16 | Enumeration Values for System-System Cache-Replacement Policy | 771 |
| M.17 | Enumeration Values for System-System Cache-Read Policy | 772 |
| M.18 | Enumeration Values for System-System Cache-Write Policy | 772 |
| M.19 | Enumeration Values for System-System Cache-Associativity | 772 |
| M.20 | Enumeration Values for System-System IRQ-Availability | 772 |
| M.21 | Enumeration Values for System-System IRQ-IRQ Trigger Type | 772 |
| M.22 | Enumeration Values for System-System DMA-Availability. | 773 |
| M.23 | Enumeration Values for Language | 773 |
| | | |
| N | Setting up Security for Server Inventory | 777 |
| | | |
| O | Documentation Updates | 779 |
| O.1 | September 07, 2007 | 779 |
| O.1.1 | Setting Up Workstation Inventory. | 779 |
| O.2 | July 27, 2007 | 780 |
| O.2.1 | Setting Up Workstation Inventory. | 780 |
| O.3 | October 19, 2006 | 780 |
| O.4 | August 26, 2006 | 780 |
| O.5 | July 14, 2006 (Support Pack 1) | 780 |
| O.6 | December 23, 2005 | 781 |
| O.7 | December 9, 2005 | 781 |
| O.8 | October 24, 2005 | 782 |
| O.9 | October 7, 2005 | 782 |

Part III Remote Management 783

19 Remote Management for NetWare Servers 785

- 19.1 Overview of RConsoleJ Components 785
 - 19.1.1 RConsoleJ Client 785
 - 19.1.2 RConsoleJ Agent 785
 - 19.1.3 RConsoleJ Proxy Agent 786
- 19.2 Setting Up RConsoleJ 786
 - 19.2.1 Loading the RConsoleJ Agent 786
 - 19.2.2 Running the RConsoleJ Client 786
 - 19.2.3 Loading the RConsoleJ Proxy Agent on a Proxy Server 787
- 19.3 RConsoleJ 787
 - 19.3.1 Scenario 1: An IP Client Controlling an IP NetWare Server 787
 - 19.3.2 Scenario 2: An IP Client Controlling an IPX NetWare Server 790
- 19.4 Loading Agents at Startup 792
 - 19.4.1 Loading RConsoleJ Agent at Startup 792
 - 19.4.2 Loading RConsoleJ Proxy Agent at Startup 792
- 19.5 Setting Up Security for RConsoleJ 792
- 19.6 Managing Remote NetWare Servers 793
 - 19.6.1 Sending Console Commands in the Server's Native Language 793
 - 19.6.2 Synchronizing RConsoleJ Client and Target NetWare Screens 794

20 Remote Management for Windows Servers 795

- 20.1 Remote Management Terminology 795
- 20.2 Understanding Remote Management for Windows Servers 795
- 20.3 Setting Up Security for Remote Management 797
 - 20.3.1 Configuring the Remote Management Policies 797
 - 20.3.2 Setting Up the Agent Password at the Managed Server 800
- 20.4 Managing Remote Windows Servers 800
 - 20.4.1 Initiating Remote Management Sessions 801
 - 20.4.2 Operating with Windows XP SP2 803
 - 20.4.3 Configuring Remote Management Ports 803
 - 20.4.4 Customizing the Permission Message 804
 - 20.4.5 Managing a Remote View Session 805
 - 20.4.6 Managing a Remote Control Session 807
 - 20.4.7 Remote Operator Identification Display 815
 - 20.4.8 Viewing the Audit Log for Remote Management Sessions 815
 - 20.4.9 Improving the Remote Management Performance 815
 - 20.4.10 Shutting Down and Restarting the Remote Management Agent 816

P Documentation Updates 819

- P.1 July 14, 2006 (Support Pack 1) 819
- P.2 December 9, 2005 819

Part IV Management and Monitoring Services 821

21 Configuring Management and Monitoring Services 823

- 21.1 Understanding Management and Monitoring Services 823
 - 21.1.1 Management Site Services 823
 - 21.1.2 Server Management 826
 - 21.1.3 Traffic Analysis 826

| | | |
|--------|--|-----|
| 21.1.4 | Novell ConsoleOne | 827 |
| 21.2 | Planning the Configuration | 828 |
| 21.2.1 | Defining Management Information Needs | 828 |
| 21.2.2 | Planning a Strategy to Manage Your Network | 828 |
| 21.3 | Role-Based Administration | 830 |
| 21.3.1 | Novell ZENworks Management Site | 830 |
| 21.3.2 | General Novell ZENworks Server Management Roles | 831 |
| 21.3.3 | Novell ZENworks Server Management Role-Based Modules and Roles | 832 |
| 21.3.4 | Configuring Role-Based Administration | 844 |
| 21.4 | Configuring Management and Monitoring Services | 846 |
| 21.4.1 | Stopping and Starting Management and Monitoring Services | 846 |
| 21.4.2 | Setting Up Discovery and Starting Back-End Processes | 847 |
| 21.4.3 | Setting Up the Alarm Management System | 848 |
| 21.4.4 | Setting Up Monitoring | 848 |
| 21.4.5 | Setting Up the Traffic Analysis Agent | 848 |

22 Using Novell ConsoleOne with Management and Monitoring Services 851

| | | |
|--------|--|-----|
| 22.1 | Navigating the Novell ZENworks Server Management Namespace | 851 |
| 22.2 | Selecting Novell ZENworks Server Management Options | 853 |
| 22.2.1 | Views | 853 |
| 22.2.2 | Properties | 854 |
| 22.2.3 | Actions | 854 |
| 22.3 | Working with Views | 854 |
| 22.3.1 | Changing the Appearance of a View | 855 |
| 22.3.2 | Modifying Columns | 856 |
| 22.3.3 | Filtering Views | 857 |
| 22.3.4 | Sorting Views | 858 |
| 22.3.5 | Printing a View | 858 |
| 22.3.6 | Exporting a View | 858 |
| 22.3.7 | Saving Views | 859 |
| 22.3.8 | Deleting and Renaming Custom Views | 859 |
| 22.3.9 | Displaying Multiple Views in Novell ConsoleOne Views | 860 |

23 Understanding Network Discovery and Atlas Management 863

| | | |
|--------|--|-----|
| 23.1 | Understanding Network Discovery | 864 |
| 23.1.1 | Discovery Components | 864 |
| 23.1.2 | Discovery Process | 874 |
| 23.1.3 | What Is Discovered | 881 |
| 23.1.4 | File-Based Discovery | 890 |
| 23.1.5 | Discovery Console | 892 |
| 23.1.6 | Effects of Discovery on Maps | 895 |
| 23.2 | Setting Up Discovery | 897 |
| 23.2.1 | Starting Discovery | 899 |
| 23.2.2 | Checking the Status of Initial Discovery | 899 |
| 23.2.3 | Checking the Results of Discovery | 900 |
| 23.2.4 | Changing the Default Configuration | 901 |
| 23.2.5 | Configuring the Java Processes | 906 |
| 23.2.6 | Unloading the Management Server | 908 |
| 23.3 | Managing the Atlas | 908 |
| 23.3.1 | Using the Atlas | 909 |
| 23.3.2 | VLAN Atlas | 915 |
| 23.3.3 | Using Unified Views | 916 |

24 Understanding Alarm Management 919

| | | |
|--------|---|-----|
| 24.1 | Understanding the Alarm Management System | 919 |
| 24.1.1 | Alarm Management System Components | 920 |
| 24.2 | Managing the Alarm Management System | 924 |
| 24.2.1 | Recognizing Alarm Indicators | 924 |
| 24.2.2 | Viewing Alarms | 925 |
| 24.2.3 | Enabling and Disabling Alarms | 929 |
| 24.2.4 | Resolving Alarms | 929 |
| 24.2.5 | Deleting Alarms | 931 |
| 24.2.6 | Performing Actions on Alarm Templates | 933 |
| 24.3 | Managing the Rule-Based Alarm Management System | 935 |
| 24.3.1 | Understanding the Properties | 936 |
| 24.3.2 | Understanding the Conditions | 936 |
| 24.3.3 | Understanding the Actions | 940 |
| 24.3.4 | Performing Actions on Rules | 947 |
| 24.4 | Maintaining the Alarm Management System | 952 |
| 24.5 | Troubleshooting the Alarm Management System | 952 |

25 Understanding Server Management 953

| | | |
|--------|--|-----|
| 25.1 | Understanding Server Management | 954 |
| 25.1.1 | SNMP-Based Server Management | 955 |
| 25.1.2 | SNMP Agent Functions | 956 |
| 25.2 | Planning for Server Management | 956 |
| 25.2.1 | Creating a Baseline of Typical Server Activity | 957 |
| 25.2.2 | Using the Baseline Document | 957 |
| 25.2.3 | Server Baseline Document Tips | 957 |
| 25.3 | Optimizing Server Management | 958 |
| 25.3.1 | Setting Default Trends and Thresholds | 959 |
| 25.3.2 | Controlling Alarm Generation | 963 |
| 25.3.3 | Defining Recipients for SNMP Alarms | 965 |
| 25.4 | Managing Servers | 966 |
| 25.4.1 | Displaying Server Configuration Information | 966 |
| 25.4.2 | Displaying Summary Data | 967 |
| 25.4.3 | Viewing Trend Data | 968 |
| 25.4.4 | Managing Trend Samplings | 970 |
| 25.4.5 | Configuring Server Parameters | 972 |
| 25.4.6 | Executing Server Commands | 972 |
| 25.4.7 | Management Site Server Status | 973 |
| 25.5 | Object Hierarchy and View Details | 974 |
| 25.5.1 | Object Hierarchy | 974 |
| 25.5.2 | Object View Details | 975 |

26 Using the MIB Tools 995

| | | |
|--------|--|------|
| 26.1 | Understanding MIB Tools | 995 |
| 26.1.1 | About MIBs | 995 |
| 26.1.2 | Understanding the SNMP MIB Compiler | 995 |
| 26.1.3 | Understanding the SNMP MIB Browser | 997 |
| 26.1.4 | Managing Devices with MIB Tools | 999 |
| 26.1.5 | Trap Definitions | 999 |
| 26.2 | Configuring MIBs and Setting Up MIB Tools | 1004 |
| 26.2.1 | Annotating Third-Party MIBs for Integration with Novell ZENworks Server Management | 1005 |
| 26.2.2 | Compiling MIBs for SNMP-Manageable Nodes | 1006 |
| 26.3 | Using the MIB Browser | 1007 |

| | | |
|--------|---|------|
| 26.3.1 | Browsing the MIB Tree | 1007 |
| 26.3.2 | Viewing the Values of an Object and Its Child Nodes | 1009 |
| 26.3.3 | Configuring a Node by Setting Object Values | 1010 |
| 26.3.4 | Modifying SNMP Preferences | 1010 |
| 26.3.5 | Modifying Instances of an SNMP Table | 1010 |
| 26.3.6 | Forming Tables of Scalar Objects | 1013 |
| 26.3.7 | Graphing SNMP Request Results | 1014 |
| 26.3.8 | Using a Profile for Tables and Graphs | 1016 |
| 26.4 | Maintaining MIBs | 1016 |

27 Using the Probe Manageability Tool 1019

| | | |
|--------|--|------|
| 27.1 | Invoking the Probe Manageability Tool | 1019 |
| 27.2 | Working with the Probe Manageability Tool | 1020 |
| 27.2.1 | Adding a Node to the List | 1020 |
| 27.2.2 | Adding Multiple Nodes at a Time to the List | 1020 |
| 27.2.3 | Deleting a Node from the List | 1021 |
| 27.2.4 | Starting the Probe Manageability Operation | 1021 |
| 27.2.5 | Stopping the Probe Manageability Operation | 1021 |
| 27.2.6 | Viewing the Probe Manageability Log for a Node | 1021 |

28 Monitoring Services 1023

| | | |
|--------|---|------|
| 28.1 | Understanding Monitoring Services | 1023 |
| 28.1.1 | Role-Based Services for Using the Monitoring Services | 1025 |
| 28.2 | Monitoring Services on Target Nodes | 1025 |
| 28.2.1 | Defining the Targets for Monitoring Services | 1025 |
| 28.2.2 | Displaying Test Results Data | 1027 |
| 28.2.3 | Changing the Test Options for a Node | 1028 |
| 28.2.4 | Adding Services for Monitoring | 1028 |

29 Understanding Traffic Analysis 1029

| | | |
|--------|---|------|
| 29.1 | Understanding Traffic Analysis | 1029 |
| 29.1.1 | Traffic Analysis Components | 1029 |
| 29.1.2 | Communication Between Traffic Analysis Components | 1030 |
| 29.1.3 | Traffic Analysis Features | 1031 |
| 29.1.4 | Traffic Analysis Fundamentals | 1032 |
| 29.2 | Planning for Segment Monitoring | 1042 |
| 29.2.1 | Creating a Baseline of Typical Segment Activity | 1042 |
| 29.2.2 | Using the Baseline Document | 1042 |
| 29.2.3 | Segment Baseline Document Tips | 1043 |
| 29.3 | Preparing to Analyze Network Traffic | 1044 |
| 29.3.1 | Selecting the Preferred RMON Agent | 1045 |
| 29.3.2 | Setting Up SNMP Parameters | 1046 |
| 29.4 | Analyzing Network Traffic | 1047 |
| 29.4.1 | Analyzing Traffic on Segments | 1047 |
| 29.4.2 | Analyzing Traffic on Nodes Connected to a Segment | 1055 |
| 29.4.3 | Capturing Packets | 1063 |
| 29.4.4 | Displaying Captured Packets | 1066 |
| 29.4.5 | Analyzing Traffic Generated by Protocols in Your Network | 1073 |
| 29.4.6 | Analyzing Traffic on Switches | 1076 |
| 29.5 | Optimizing Traffic Analysis | 1078 |
| 29.5.1 | Choosing Options to Display Stations on a Segment | 1079 |
| 29.5.2 | Choosing Options to Display Trend Statistics | 1080 |
| 29.5.3 | Choosing Options to Display the Top Nodes Graph | 1083 |
| 29.5.4 | Choosing Statistics to Display in the Unified Port Traffic View | 1084 |

| | | |
|--------|---|------|
| 29.5.5 | Choosing Options to Display a Captured Packet | 1085 |
| 29.5.6 | Configuring Alarm Options from the Set Alarm Dialog Box | 1085 |
| 29.5.7 | Configuring the Monitor Nodes for Inactivity View | 1088 |
| 29.6 | Understanding the Traffic Analysis Agents | 1089 |
| 29.7 | Using the Traffic Analysis Agent for NetWare | 1090 |
| 29.7.1 | Planning to Install the Traffic Analysis Agent for NetWare | 1091 |
| 29.7.2 | Optimizing the Traffic Analysis Agent for NetWare Performance | 1092 |
| 29.7.3 | Using the Console Utility of the Traffic Analysis Agent for NetWare | 1098 |
| 29.8 | Using the Traffic Analysis Agent for Windows | 1104 |
| 29.8.1 | Changes Made During Installation | 1105 |
| 29.8.2 | Planning to Install the Traffic Analysis Agent for Windows | 1106 |
| 29.8.3 | Optimizing the Traffic Analysis Agent for Windows | 1109 |
| 29.8.4 | Using LANZCON | 1112 |

30 Customizing the Agent Configuration 1115

| | | |
|--------|--|------|
| 30.1 | Agent Files | 1115 |
| 30.1.1 | Management Agent for NetWare Files | 1115 |
| 30.1.2 | Management Agent for Windows Server Files | 1117 |
| 30.2 | Customizing the Management Agent for NetWare | 1118 |
| 30.2.1 | servinst.nlm Load Parameters | 1118 |
| 30.2.2 | hostmib.nlm Load Parameters | 1119 |
| 30.2.3 | ntrend.nlm Load Parameters | 1120 |
| 30.3 | Customizing the Management Agent for Windows Server | 1120 |
| 30.3.1 | Configuring the Management Agent for Windows Server | 1121 |
| 30.3.2 | Collecting Events from Custom Event Log Types | 1121 |
| 30.3.3 | Specifying Negative Filter Conditions in the Nttrap.ini File | 1122 |
| 30.4 | Third-Party Agent Configuration | 1122 |
| 30.4.1 | Ensuring that Traps Are Received | 1122 |
| 30.4.2 | Integrating Vendor-Specific SNMP Traps | 1122 |
| 30.5 | Advanced Trending Agent | 1123 |
| 30.5.1 | What Is the Advanced Trending Agent? | 1123 |
| 30.5.2 | Configuring the Trend Variables | 1123 |
| 30.5.3 | Configuring the Advanced Trending Agent on Linux | 1124 |
| 30.5.4 | Configuring the Advanced Trending Agent on All Platforms | 1125 |
| 30.5.5 | Quick Reference Table | 1126 |
| 30.5.6 | Refreshing Configuration Settings | 1127 |
| 30.5.7 | Installing the Advanced Trending Agent | 1128 |
| 30.6 | Management and Monitoring Services for Linux | 1129 |
| 30.6.1 | Providing Real Time Statistical Information | 1129 |
| 30.6.2 | Generating Traps for System Events | 1130 |
| 30.6.3 | Providing History Collection Information | 1132 |
| 30.6.4 | Linux Management Views | 1132 |
| 30.6.5 | Linux Server Health Reports | 1133 |

31 Protocol Decodes Suites Supported by Novell ZENworks Server Management 1135

| | | |
|------|---|------|
| 31.1 | Novell NetWare Protocol Suite | 1135 |
| 31.2 | Network File System Protocol Suite | 1137 |
| 31.3 | Systems Network Architecture Protocol Suite | 1137 |
| 31.4 | AppleTalk Protocol Suite | 1138 |
| 31.5 | TCP/IP Protocol Suite | 1139 |

| | |
|---|-------------|
| 32 Novell ZENworks Management and Monitoring Services Database | 1143 |
| 32.1 Understanding the Novell ZENworks Server Management Database | 1143 |
| 32.1.1 Running the Database | 1143 |
| 32.1.2 Database Caching | 1144 |
| 32.2 Backing Up the Topology/Alarm Database. | 1144 |
| 32.3 Changing Database Passwords | 1144 |
| 32.4 Emptying the Database | 1144 |
| | |
| 33 Using Reports in Management and Monitoring Services | 1145 |
| 33.1 Understanding Management and Monitoring Services Reports | 1145 |
| 33.1.1 About the Topology Reports | 1145 |
| 33.1.2 About the Alarm Reports | 1148 |
| 33.1.3 About the Health Reports. | 1149 |
| 33.2 Managing Reporting | 1151 |
| 33.2.1 Managing the Topology Reports | 1151 |
| 33.2.2 Managing the Server Management Health Reports. | 1151 |
| | |
| 34 Using SNMP Community Strings | 1159 |
| 34.1 About SNMP Community Strings | 1159 |
| 34.1.1 SNMP Security. | 1159 |
| 34.2 Setting the SNMP Community Strings | 1160 |
| 34.2.1 Setting the SNMP Community String: Novell NetWare Server | 1160 |
| 34.2.2 Setting the SNMP Community String: Novell ConsoleOne | 1162 |
| 34.2.3 Setting Community Strings for an Individual Node. | 1162 |
| 34.2.4 Setting the SNMP Community String: Windows | 1163 |
| | |
| 35 Understanding the View Builder | 1165 |
| 35.1 Creating a Name-Value Pairs View Component | 1167 |
| 35.1.1 Adding a Name-Value Pair. | 1168 |
| 35.1.2 Editing the Name-Value Pairs View Component | 1168 |
| 35.1.3 Deleting the Name-Value Pairs View | 1168 |
| 35.2 Creating an Alarm View Component | 1169 |
| 35.2.1 Editing the Alarms View Component | 1169 |
| 35.2.2 Deleting the Alarm View Component | 1169 |
| 35.3 Creating a Table View Component. | 1170 |
| 35.3.1 Editing an Table View | 1171 |
| 35.3.2 Deleting a Table View | 1171 |
| 35.4 Creating a Graph View Component | 1171 |
| 35.4.1 Adding Graph Details. | 1172 |
| 35.4.2 Editing the Graph View Component. | 1172 |
| 35.4.3 Deleting the Graph View Component. | 1172 |
| 35.5 Setting the Criteria for the View to Appear | 1172 |
| | |
| 36 Understanding Trap Configuration | 1175 |
| 36.1 Understanding Trap Configuration | 1175 |
| 36.1.1 The Configuration Agents | 1175 |
| 36.1.2 Trap Configuration Management Console | 1175 |
| 36.2 Configuring Traps Using Trap Configuration Page. | 1176 |
| 36.2.1 Enabling and Disabling the Traps | 1176 |
| 36.2.2 Changing the Interval of a Trap | 1177 |
| 36.2.3 Selecting NetWare Servers to Apply a Trap Configuration | 1177 |

| | | |
|--------|---|------|
| 36.2.4 | Viewing the Trap Configuration Status | 1178 |
| 36.2.5 | Using the Command Line Option | 1178 |
| 36.3 | Additional Trap Configuration Features | 1179 |
| 36.3.1 | Filtering the Traps | 1179 |
| 36.3.2 | Sorting the Traps | 1180 |
| 36.3.3 | Managing Profiles | 1180 |
| 36.3.4 | Viewing Current Configuration of a Server | 1181 |
| 36.3.5 | A Use Case: Configuring NetWare Servers | 1181 |

Q Setting up Security for Management and Monitoring Services 1183

R Documentation Updates 1185

| | | |
|-----|--------------------------------------|------|
| R.1 | August 16, 2006. | 1185 |
| R.2 | July 14, 2006 (Support Pack 1) | 1185 |
| R.3 | December 9, 2005 | 1186 |

About This Guide

This guide describes how to administer Novell® ZENworks® 7 Server Management with Support Pack 1 (SP1). The guide is divided into the following sections:

- ◆ Part I, “Policy and Distribution Services,” on page 27
- ◆ Part II, “Server Inventory,” on page 443
- ◆ Part III, “Remote Management,” on page 783
- ◆ Part IV, “Management and Monitoring Services,” on page 821

Audience

This guide is intended for ZENworks administrators.

Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comments feature at the bottom of each page of the online documentation, or go to www.novell.com/documentation/feedback.html and enter your comments there.

Documentation Updates

For the most recent version of this guide, the Web HTML and updated PDF versions are available on the [ZENworks 7 Web site \(http://www.novell.com/documentation/beta/zenworks7/index.html\)](http://www.novell.com/documentation/beta/zenworks7/index.html).

Additional Documentation

For the latest documentation on installing ZENworks 7 Server Management with SP1, see the *Novell ZENworks 7 Server Management Installation Guide*.

Documentation Conventions

In Novell documentation, a greater-than symbol (>) is used to separate actions within a step and items in a cross-reference path.

A trademark symbol (® , ™ , etc.) denotes a Novell trademark. An asterisk (*) denotes a third-party trademark.

When a single pathname can be written with a backslash for some platforms or a forward slash for other platforms, the pathname is presented with a backslash. Users of platforms that require a forward slash, such as Linux or UNIX*, should use forward slashes as required by your software.

Policy and Distribution Services

Novell® ZENworks® 7 Server Management Policy and Distribution Services is a software, configuration, and behavioral management system for servers. Through Policy and Distribution Services, you can:

- ◆ Control the versions of software installed on servers throughout your network
- ◆ Define and enforce a standard configuration on any given set of servers
- ◆ Control the behavior of servers in given situations, such as downing a server, backing up volumes, managing thresholds exceeded, and so on

Policy and Distribution Services has three components:

- ◆ **Tiered Electronic Distribution:** Simplifies data delivery and server policy implementation
- ◆ **Server Policies:** Simplifies configuration and management of your servers
- ◆ **Server Software Packages:** Simplifies the installation of software

You can administer Policy and Distribution Services by using the following:

- ◆ **ConsoleOne 1.3.6**, where you can create and configure Server Management objects and perform management tasks for Policy and Distribution Services.
- ◆ **ZENworks Server Management role in Novell iManager**, where you can perform management tasks for Policy and Distribution Services using iManager from any workstation where Internet Explorer 5.5 or later is available.

This Policy and Distribution Services documentation contains the following sections:

- ◆ [Chapter 1, “Post-Installation Setup,” on page 29](#) (After installing ZENworks 7 Server Management for the first time, use this section to complete a full configuration of your policies and your distribution system.)
- ◆ [Chapter 2, “Novell iManager,” on page 63](#)
- ◆ [Chapter 3, “Tiered Electronic Distribution,” on page 85](#)
- ◆ [Chapter 4, “Server Policies,” on page 193](#)
- ◆ [Chapter 5, “Server Software Packages,” on page 237](#)
- ◆ [Chapter 6, “Desktop Application Distribution,” on page 273](#)
- ◆ [Chapter 7, “Security in Policy and Distribution Services,” on page 299](#)
- ◆ [Chapter 8, “Scheduling,” on page 317](#)
- ◆ [Chapter 9, “Variables,” on page 341](#)
- ◆ [Chapter 10, “ZENworks Database,” on page 351](#)
- ◆ [Chapter 11, “Reporting,” on page 363](#)
- ◆ [Appendix A, “Distribution Types,” on page 383](#)
- ◆ [Appendix B, “Schedule Types,” on page 399](#)
- ◆ [Appendix C, “Server Console Commands,” on page 405](#)
- ◆ [Appendix D, “Load/Unload Actions,” on page 411](#)

- ◆ [Appendix E, “Requirements for Server Software Packages,” on page 413](#)
- ◆ [Appendix F, “Registry Entries for Server Software Package Components,” on page 421](#)
- ◆ [Appendix G, “Client Access in Linux,” on page 425](#)
- ◆ [Appendix H, “Configuration Planning Worksheet,” on page 427](#)
- ◆ [Appendix I, “Documentation Updates,” on page 437](#)

Post-Installation Setup

1

To use the Tiered Electronic Distribution capability of Novell® ZENworks® Server Management effectively, you must correctly install and configure its components on your network. You should have already performed a basic installation of Policy and Distribution Services (see “[Installation on NetWare and Windows Servers](#)” in the *Novell ZENworks 7 Server Management Installation Guide*).

For information on configuring policies, see [Chapter 4, “Server Policies,” on page 193](#).

This section provides you with the concepts, a [planning worksheet](#), and instructions to help you further configure your Tiered Electronic Distribution system. For more detailed information, see [Chapter 3, “Tiered Electronic Distribution,” on page 85](#).

The information provided in the following sections will help you to add new Distributors as needed, finish installing the Subscriber software as needed, configure a Distributor’s routing hierarchy, create some Distributions, and send those Distributions:

- ◆ [Section 1.1, “Planning Your Distribution System,” on page 29](#)

In this section, you can use the [planning worksheet](#) to keep track of the decisions you need to make. Then you can easily perform your planned configurations from the information on the planning worksheet.

- ◆ [Section 1.2, “Configuring Your Distribution System,” on page 50](#)

This section provides the steps for configuring your distribution system.

- ◆ [Section 1.3, “Managing Your Distribution System,” on page 60](#)

This section provides an overview of how you can manage your distribution system using Novell ConsoleOne® and Novell iManager.

- ◆ [Appendix H, “Configuration Planning Worksheet,” on page 427](#)

The planning worksheet contains basic information for each worksheet entry. It also contains links to where you can view more information to better understand a worksheet entry.

The worksheet should not be used in place of the procedures in [Section 1.2, “Configuring Your Distribution System,” on page 50](#), because the worksheet contains only planning information; it does not contain information for the procedures that are not planned.

1.1 Planning Your Distribution System

Use these sections in the following order:

1. [“Overview” on page 30](#)
2. [“Selecting Your Distributions” on page 32](#)
3. [“Understanding Your Network Topology” on page 36](#)
4. [“Are Additional Distributors Needed?” on page 37](#)
5. [“Other Subscribers To Be Installed?” on page 41](#)
6. [“Determining the Distribution Flow” on page 41](#)
7. [“Understanding Distribution Security” on page 44](#)

8. [“Determining the Channels for the Distributions” on page 46](#)
9. [“Determining Subscribers’ Subscriptions” on page 47](#)
10. [“Determining the Distribution Schedules” on page 48](#)

1.1.1 Overview

Policy and Distribution Services contains three components:

- ◆ **Tiered Electronic Distribution** is a distribution system for your network.
 - ◆ It is a way to manage your network servers through the distribution of electronic data between servers.
 - ◆ It uses a tiered architecture for distribution efficiency. For example, workload sharing: one server can service many others, then each of those many servers can also service many more, and so on to any number of tiers.
 - ◆ It provides Distribution scheduling for efficient bandwidth usage, such as distributing during off-peak hours.
 - ◆ It provides security to prevent unauthorized tampering with the Distributions.
- ◆ **Server Policies** is a system for managing the configuration and behavior of your servers.
- ◆ **Server Software Packages** is a feature for automating the installation and upgrading of software on your servers.

Tiered Electronic Distribution is usually involved when you use any of these components, because most policies and all Server Software Packages are distributed. Therefore, in this section we will concentrate on understanding and configuring Tiered Electronic Distribution. See the following sections for more information on the other two components of Policy and Distribution Services:

- ◆ [Chapter 4, “Server Policies,” on page 193](#)
- ◆ [Chapter 5, “Server Software Packages,” on page 237](#)

The following sections provide basic information that will help you to understand Tiered Electronic Distribution and what you will need to know to configure it:

- ◆ [“What Can You Distribute?” on page 30](#)
- ◆ [“How Is Data Distributed?” on page 31](#)
- ◆ [“What Do You Need to Know to Plan Your Distribution System?” on page 31](#)

What Can You Distribute?

The types of electronic data you can distribute using Tiered Electronic Distribution include:

Table 1-1 *Distribution Types*

| Distribution Type | Content Distributed |
|---------------------|--|
| Desktop Application | Desktop Application objects and files created in ZENworks Desktop Management |
| File | Files and directories contained on the Distributor server’s file system |

| Distribution Type | Content Distributed |
|-------------------|---|
| FTP | Files and directories from an FTP source |
| HTTP | Content from an HTTP source |
| MSI | Contains software to be installed in a Windows* environment by the MSI engine |
| Policy Package | Policies for controlling servers |
| RPM | RPM packages for Linux and Solaris* servers (but only for Solaris if RPM is installed to the Solaris machine) |
| Software Package | Server Software Packages for automatically installing or upgrading software on your servers |

From this list, you can see that there is a variety of electronic data types that you can distribute to your servers.

How Is Data Distributed?

Tiered Electronic Distribution sends Distribution files from Distributor servers to Subscriber servers. The basic distribution process is as follows:

1. Decide what you want to distribute.
2. Create the Distribution.
3. Create a Channel for the Distribution.
4. Determine which Subscriber servers need this Distribution.
5. Subscribe the Subscriber servers to the Distribution's Channel.
6. Make sure the applicable schedules are set (Build, Send, and Extract).
7. Send the Distribution by refreshing the Distributor, which causes the Distribution to be built according to the Distribution's Build schedule, and sent according to the Channel's Send schedule.
8. The Distribution is extracted on the Subscriber servers according to their Extract schedules.
9. The Distributions are used by the Subscriber servers according to the Distribution's type.

From this process, you can see that there are several components of Tiered Electronic Distribution that will need to be created and configured. For more information, see [Section 3.2.2, "The Basic Distribution Process," on page 88](#) and [Section 3.10.1, "Understanding the Distribution Processes," on page 172](#).

What Do You Need to Know to Plan Your Distribution System?

- ◆ The Distributions that you want, including:
 - ◆ Whether you want to distribute server files, HTTP content, FTP content, or RPM packages
 - ◆ If there are any desktop applications to be distributed (affects how you set up Subscriber objects when you have multiple trees)
 - ◆ Which policies you needed for managing your servers
 - ◆ What server software should have automated installation
- ◆ Whether you'll need additional Distributors

- ◆ Whether you have both Novell eDirectory™ 8.7.3 and NDS® 6.x or 7.x in your environment, which adversely affects Distributors (a workaround is available)
- ◆ How many databases you'll need for reporting purposes
- ◆ Whether you need to complete installation of the Subscriber software to your servers
- ◆ Which Subscribers need which Distributions
- ◆ Your network's topology (server platforms, slow WANs, firewalls, Network Address Translation [NAT], multiple trees, and so on)
- ◆ The system resource and server behavior issues that Tiered Electronic Distribution might create
- ◆ Whether you need to encrypt Distributions for certain servers
- ◆ Whether you can use Subscriber Groups for channeling Distributions
- ◆ How you want the Distributions to flow to the Subscriber servers (the tiered distribution model)
- ◆ How you want to schedule the distribution processes to minimize network traffic during business hours

To determine the above information, continue with [Section 1.1.2, "Selecting Your Distributions," on page 32](#).

1.1.2 Selecting Your Distributions

This section provides you with basic information for each Distribution type.

You can build your distribution system incrementally by adding Distributions a few at a time, then adding Distributors when needed. You can revisit this process at any time to add new Distributions.

Print a copy of the [Appendix H, "Configuration Planning Worksheet," on page 427](#). Worksheet fill-in instructions are given as you review the planning sections.

Review the following Distribution type sections to select which ones you want to create at this time. [Planning worksheet](#) entries are provided for each Distribution type.

- ◆ ["Desktop Application" on page 32](#)
- ◆ ["File" on page 33](#)
- ◆ ["FTP" on page 33](#)
- ◆ ["HTTP" on page 34](#)
- ◆ ["MSI" on page 34](#)
- ◆ ["Policy Package" on page 34](#)
- ◆ ["RPM" on page 36](#)
- ◆ ["Software Package" on page 36](#)

Desktop Application

This Distribution type allows you to distribute Application objects and associated files to specified locations on the eDirectory tree and target Subscriber servers.

For information on integration with Desktop Management, see [Chapter 6, "Desktop Application Distribution," on page 273](#).

For information on the Desktop Application type of Distribution, see [“Desktop Application” on page 117](#).

Determine whether you want to create a Desktop Application Distribution at this time:

CONFIGURATION PLANNING WORKSHEET

If you want to maintain trustee rights in the Distribution, under [item 3](#) and [item 20](#), indicate that you have Desktop Application Distributions, and therefore each server that receives Desktop Application Distributions must have its Subscriber object and NCP™ server object in the same tree.

Under [item 19](#), enter Desktop Application as the type of Distribution to be created. Also indicate the following:

- ◆ A name for the Distribution that indicates its purpose
 - ◆ Names of the servers that need a Desktop Application Distribution
-

File

With this Distribution type you can select files and/or directories from the Distributor server’s file system to distribute to a selected location on the Subscriber server’s file system.

A Distribution Wizard is available for automating the process of creating the File and FTP types of Distributions. For more information, see [Section 3.4.12, “Using the Distribution Wizard,” on page 143](#).

For information on the File type of Distribution, see [“File” on page 118](#).

Determine whether you want to create a File Distribution at this time:

CONFIGURATION PLANNING WORKSHEET

Under [item 19](#), enter File as the type of Distribution to be created. Also indicate the following:

- ◆ A name for the Distribution that indicates its purpose
 - ◆ Names of the servers that need a File Distribution
-

FTP

With this Distribution type you can create a Distribution consisting of files from one or more FTP sources. Each source can contain one or more directories and/or files.

A Distribution Wizard is available for automating the process of creating the File and FTP types of Distributions. For more information, see [Section 3.4.12, “Using the Distribution Wizard,” on page 143](#).

For information on the FTP type of Distribution, see [“FTP” on page 119](#).

Determine whether you want to create an FTP Distribution at this time:

CONFIGURATION PLANNING WORKSHEET

Under **item 19**, enter FTP as the type of Distribution to be created. Also indicate the following:

- ◆ A name for the Distribution that indicates its purpose
 - ◆ Names of the servers that need an FTP Distribution
-

HTTP

With this Distribution type you can create a Distribution consisting of one or more HTTP sources. Each source can contain one or more target entries.

For information on the HTTP type of Distribution, see **“HTTP” on page 120**.

Determine whether you want to create an HTTP Distribution at this time:

CONFIGURATION PLANNING WORKSHEET

Under **item 19**, enter HTTP as the type of Distribution to be created. Also indicate the following:

- ◆ A name for the Distribution that indicates its purpose
 - ◆ Names of the servers that need an HTTP Distribution
-

MSI

This is a Distribution of MSI packages that are installed by the MSI engine in a Windows environment.

For information on the MSI type of Distributions, see **“MSI” on page 120**.

Determine whether you want to create an MSI Distribution at this time:

CONFIGURATION PLANNING WORKSHEET

Under **item 19**, enter MSI as the type of Distribution to be created. Also indicate the following:

- ◆ A name for the Distribution that indicates its purpose
 - ◆ Names of the servers that need an MSI Distribution
-

Policy Package

This Distribution type provides the mechanism for applying any of the policies in **Table 1-2** to Subscriber servers:

Table 1-2 Policies

| Policy | Description |
|------------|---|
| Copy Files | Enables copying of files on a server from one location to another by using policy configurations. |

| Policy | Description |
|--------------------------------|---|
| NetWare Set Parameters | Specifies and optimizes selected NetWare® Set Parameters for a server or group of servers. |
| Prohibited File | Used to monitor and enforce the deletion or moving of unauthorized files from a specified volume/drive or directory/folder. |
| Scheduled Down | Schedules when a server should go down, and whether it should be brought back up automatically. |
| Scheduled Load/Unload | Automates the loading and unloading order of NLM™ and Java* Class processes for the selected servers, and for starting and stopping Windows services. |
| Search | Used in Server Management to enable the Distributor Agent to locate and use policies in the Service Location Package. |
| Server Down Process | Controls which processes to follow and which conditions to meet before downing a server. |
| Server Scripts | Automates script usage on your servers. |
| SMTP Host | Sets the TCP/IP address of the relay host that processes outbound Internet e-mail. |
| SNMP Community Strings | Allows you to receive and respond to SNMP requests. |
| SNMP Trap Targets | Sets SNMP trap targets for associated eDirectory objects for reporting purposes. |
| Text File Changes | Automates changes to text files. |
| Tiered Electronic Distribution | Sets defaults for the Distributor and Subscriber objects. |
| ZENworks Database | <p>Sets the DN for locating a ZENworks Database object and the path to the database file. The database is used by Policy and Distribution Services for logging successes and failures that are used in creating reports.</p> <p>The database location specified during installation can be overridden by creating and enabling this policy.</p> |
| ZENworks Server Management | Contains basic configuration parameters for Policy and Distribution Services, such as status logging, defining the server console prompt for the Policy/Package Agent, setting its working path, and setting a database purging limit. |

For more information on each policy, see [Section 4.1.6, “Server Policy Descriptions,”](#) on page 200.

For information on policies and policy packages, see [Chapter 4, “Server Policies,”](#) on page 193.

For more information on the Policy Package type of Distribution, see [“Policy Package”](#) on page 121.

Determine whether you want to create a Policy Package Distribution at this time:

CONFIGURATION PLANNING WORKSHEET

Under **item 19**, enter Policy Package as the type of Distribution to be created. Also indicate the following:

- ◆ Names of the policies
 - ◆ For each policy, names of servers that need the policy
-

RPM

This is a Linux or Solaris platform Distribution. You can distribute Red Hat* Package Manager (RPM) packages using the RPM Distribution.

For information on the RPM type of Distribution, see [“RPM” on page 121](#).

Determine whether you want to create an RPM Distribution at this time:

CONFIGURATION PLANNING WORKSHEET

Under **item 19**, enter RPM as the type of Distribution to be created. Also indicate the following:

- ◆ A name for the Distribution that indicates its purpose
 - ◆ Names of the servers that need an RPM Distribution
-

Software Package

This Distribution type allows you to distribute Server Software Packages that you create in ConsoleOne in the Server Software Package namespace. You first create a `.spk` file, then compile it into the `.cpk` file that is distributed.

For information on Server Software Packages, see [Chapter 5, “Server Software Packages,” on page 237](#).

For information on the Software Package Distribution type, see [“Software Package” on page 122](#).

Determine the software packages you want to create at this time:

CONFIGURATION PLANNING WORKSHEET

Under **item 19**, enter Software Package as the type of Distribution to be created. Also indicate the following:

- ◆ A name for the Distribution that indicates its purpose
 - ◆ Names of servers that need a Software Package Distribution
-

1.1.3 Understanding Your Network Topology

In order for you to efficiently manage your distribution system, you need to know your network’s topology. For example:

- ◆ What are your server platforms?
- ◆ How many servers do you have per platform?
- ◆ Where are your servers located in relation to WAN links and firewalls?

- ◆ Is Network Address Translation (NAT) being used?
- ◆ Where are your slow network links?

This type of information is used to help you configure the best distribution management solution for your network.

To obtain information concerning your network:

- 1 Note the trees where you extended the schema for Server Management.

CONFIGURATION PLANNING WORKSHEET

Under **item 1**, provide the names of the trees in your network where you extended the schema for Server Management.

- 2 Draw a diagram of your network structure.

You will use this diagram later to determine distribution routes.

Indicate the following on your diagram:

- ◆ Where slow links exist
- ◆ The number of servers on each LAN
- ◆ The number of servers outside a firewall
- ◆ The number of servers using NAT

- 3 Draw tree diagrams that show how your trees are currently organized. Include the main containers, such as:

- ◆ The containers that represent geographic locations (a physical tree design)
- ◆ The containers that represent the corporate organization (a logical tree design)
- ◆ The containers where servers reside (for Distributors and Subscribers)

- 4 Indicate the following on your tree diagrams:

- ◆ Where servers are located that could be Distributors (NetWare, Windows, Linux, or Solaris servers that exceed the minimum Server Management requirements)
- ◆ Containers where there are slow network connections

This should match where you indicated slow connections on your network diagram.

- 5 Indicate the following on your network diagram:

- ◆ Where the servers are located (as you just noted on the tree diagrams) that could be Distributors

1.1.4 Are Additional Distributors Needed?

When installing Policy and Distribution Services for the first time, you installed one Distributor with a database file. Generally, you'll need Distributors according to your corporate structure or geographic locations.

Distributor server workload, including the ability to complete Distribution building tasks, should also determine how many Distributors you need. For example, if you have a very large Distribution that you want built during off-peak hours, which does not need to be sent immediately, and also have virus pattern Distributions that do need to be sent immediately, you might need two different

Distributors, one with a daily refresh schedule (because you are only going to be building the Distribution once per day), and another with a frequent refresh schedule for discovering new virus pattern changes, so that their Distributions can be built and sent on time.

Use your diagrams to determine whether you need to install additional Distributors.

CONFIGURATION PLANNING WORKSHEET

Under [item 2](#), provide the names of the servers where you want to install the Distributor software.

You can always add Distributors later after you've seen how your Distributor servers handle their Distribution building and sending workload, you can determine whether to add additional Distributors for spreading that workload.

You also need to determine the following information for each Distributor:

- ◆ [“Distributor Properties” on page 38](#)
- ◆ [“Software Installation Paths” on page 39](#)
- ◆ [“Whether a Distributor Server Will Host a Server Management Database” on page 39](#)
- ◆ [“Whether Distributors Might Exist in a Mixed eDirectory Environment” on page 40](#)

Distributor Properties

You can change the following Distributor properties from the defaults during installation:

- ◆ **Object name:** If you want to rename the Distributor object, we recommend that you maintain the server's identity in the name, including the fact that it is a Distributor.
- ◆ **Container:** Plan on using the container where you previously installed Distributor objects.
If eDirectory is not installed on the Windows 2000/2003 server that you want to be a Distributor, a default container object is not displayed for that server during installation. Therefore, determine the container for that Distributor object.
- ◆ **Working directory:** You can use a different volume, drive, or directory path for the Distributor's working files than the default path.

Because the working directory has the potential to be quite large (depending on the size of the Distributions), make sure you have enough disk space.

The default volume on a NetWare server is sys:. For NetWare servers we strongly recommend that you specify a different volume.

The default working directory path for NetWare and Windows servers is:

```
\zenworks\pds\ted\dist
```

For Linux or Solaris servers the path is:

```
/var/opt/novell/zenworks/zfs/pds/ted/dist
```

The Distributor's working directory is also used whenever a Distribution is created. A directory is created under the working directory using the DN of the Distribution object.

For more information on the working directory, see [Section 3.12, “Working Directories,” on page 186](#).

CONFIGURATION PLANNING WORKSHEET

Under **item 7**, provide the property information for the Distributor that you want to be different than the defaults. This includes object names, containers for the object, and working directories.

Software Installation Paths

Server Management uses the following default installation paths:

- ◆ **NetWare:** `sys :`
You can select a different volume.
- ◆ **Windows:** `C :`
You can select a different drive.

The Linux or Solaris path cannot be changed.

CONFIGURATION PLANNING WORKSHEET

Under **item 5**, provide the installation path information for the Distributor if it is different from the default path. Include the identities of the Distributors where you have different Distributor installation paths.

Under **item 6**, provide the installation path information for the Subscriber if it is different from the default path. Include the identities of the Subscribers where you have different Subscriber installation paths.

Whether a Distributor Server Will Host a Server Management Database

You can have multiple Server Management databases in the tree, and you can install the database to both NetWare and Windows servers.

The database is used by Policy and Distribution Services to log successes and failures for the Server Policies or Tiered Electronic Distribution components. Policy and Distribution Services can function normally without a database, because it uses the `zfslog.db` file to only log information for reports. `zfslog.db` for Policy and Distribution Services does not contain any configuration information.

To determine whether you want each Distributor to have its own database, or have all Distributors share the same database, you need to determine how you want information reported. Consider the following to determine how many databases to have in the tree:

- ◆ **WAN traffic:** Tiered Electronic Distribution does not perform a large number of database updates, so the actual impact on system resources should be minimal. The greatest impact could be the time it takes to perform the transaction. However, if you have slow WAN connections, you might not want database logging to occur over the WAN.
- ◆ **Multiple Distributors:** If you have multiple Distributors in the tree, you can have one database for each, or have them share one or more databases. The type of Distributor reporting you want should determine whether to have a separate database for each. For example, are your Distributors specialized in the types of Distributions they'll send?
- ◆ **Consolidated reporting:** To have only one report for all of your Tiered Electronic Distribution information, install only one database object and file and have all Tiered Electronic Distribution Distributors log to that one file, regardless of WAN traffic considerations. Use the

ZENworks Database policy (Service Location Package) to direct all Distributors to that database file.

- ◆ **Specialized reporting:** You might want reports that are specific to a region or group of servers. You can install a database object and file for each region and have the Distributors in those regions or server groups log to that database. Use a separate ZENworks Database policy (Service Location Package) to direct each Distributor to its desired database file.

For more information, see [Chapter 10, “ZENworks Database,” on page 351](#).

IMPORTANT: Make sure you select a server for the database where you are installing the Subscriber/Policies option. The Purge Database option in the ZENworks Server Management policy (Distributed Server Package) works only if the Policy/Package Agent software and the `zfslog.db` file are located on the same server.

CONFIGURATION PLANNING WORKSHEET

Provide the following information for each Database object to be created:

- ◆ Under [item 4](#), provide the name of the Distributor server that hosts the Server Management database file.
 - ◆ Under [item 9](#), provide the installation path information that is different from the default path.
 - ◆ Under [item 10](#), provide a name for the Database object, if different from the default.
 - ◆ Under [item 11](#), provide the eDirectory container where the Database object should be created.
-

Whether Distributors Might Exist in a Mixed eDirectory Environment

Server Management can run in a mixed eDirectory environment. For example, your network might have both eDirectory 8.x and NDS[®] 6.x or 7.x installed.

However, eDirectory 8.x (only 8.6.2, 8.7.1, or 8.7.3 or later) is required for Server Management so that its objects can be placed in the tree during installation of the product. eDirectory must be installed with the master replica somewhere in your network, but not necessarily on a server where you are installing the Server Management software.

Also, ZENworks 7 Distributor servers must be running eDirectory 8.x.

The only requirement for any Server Management server is that it can communicate with the server where the eDirectory master replica (of the partition where its NCP Server object resides) is installed. Therefore, you do not need to install eDirectory on each server where you will install Server Management.

Select an IP address of any server in your tree that is using eDirectory 8.x. This can even be the IP address of the Distributor server itself, if the server is running eDirectory 8.x.

CONFIGURATION PLANNING WORKSHEET

Under [item 12](#), provide the IP address of a server using eDirectory 8.x.

1.1.5 Other Subscribers To Be Installed?

When you first installed Policy and Distribution Services, you might not have installed the software to all of your servers. If you determined that you wanted to install the Subscriber software incrementally to your servers, you can complete another stage at this time.

You can change the following Subscriber properties from the defaults during installation:

- ♦ **Object name:** If you want to rename the Subscriber object, we recommend that you maintain the server's identity in the name, including the fact that it is a Subscriber.
- ♦ **Container:** Plan on using the container where you previously installed Subscriber objects. You should place Subscriber server objects in containers matching their operating systems. For example, a NetWare container for NetWare servers, and a Windows container for Windows servers.

If eDirectory is not installed on the Windows 2000/2003 server that you want to be a Subscriber, a default container object is not displayed for that server during installation. Therefore, determine the container for that Subscriber object.

- ♦ **Working directory:** You can use a different volume, drive, or directory path for the Subscriber's working files than the default path.

Because the working directory has the potential to be quite large (depending on the size of the Distributions), make sure you have enough disk space. The default volume on a NetWare server is sys:. For NetWare servers we strongly recommend that you specify a different volume.

You might need to provide different paths for your Subscriber servers. For example, sys: for NetWare servers and D: for Windows servers. You can use variables for path data, such as the volume/drive designation. For more information, see [Chapter 9, "Variables," on page 341](#).

The default working directory path for NetWare and Windows servers is:

```
\zenworks\pds\ted\sub
```

For Linux and Solaris servers, the path is:

```
/var/opt/novell/zenworks/zfs/pds/ted/sub
```

For more information on working directories, see [Section 3.12, "Working Directories," on page 186](#).

CONFIGURATION PLANNING WORKSHEET

Under [item 3](#), provide the names of the servers where you want to install the Subscriber software at this time.

For each Subscriber to be installed, under [item 8](#), provide the property information that you want to be different than the defaults. This includes object names, containers for the object, and working directories.

1.1.6 Determining the Distribution Flow

The following sections provide information for determining distribution routes:

- ♦ ["Understanding Distribution Routes" on page 42](#)
- ♦ ["Selecting Subscribers for the Distribution Routes" on page 43](#)
- ♦ ["Configuring the Distribution Routes" on page 43](#)

For more detailed information, see [Section 3.3.2, “Understanding Distribution Routing,”](#) on [page 97](#).

Understanding Distribution Routes

Each Distributor has a routing hierarchy that provides it with a hierarchical path for sending its Distributions. The routing hierarchy contains a list of Subscribers. The hierarchy of Subscribers can be many levels deep.

Subscribers in a Distributor’s routing hierarchy do not need to also be recipients of the Distributions from that Distributor. A Subscriber can merely act as a proxy for the Distributor to pass Distributions to other Subscribers.

Not all Subscribers are needed in a routing hierarchy; only the ones used to pass Distributions on to other Subscriber servers. Most of your network’s Subscriber servers will likely be end-node Subscribers; meaning, Subscribers that only receive and extract the Distributions.

The Distributor determines the most efficient route to any given Subscriber as follows:

1. The Distributor identifies the Subscriber that is to receive the Distribution.
2. The Distributor determines whether that Subscriber has a parent Subscriber.
3. If the Subscriber has a parent Subscriber, the Distributor checks its routing hierarchy for that parent Subscriber:
 - a. If the parent Subscriber is in the routing hierarchy, the Distributor uses that route to send the Distribution to the Subscriber.
 - b. If the parent Subscriber is not in the routing hierarchy, the Distributor sends the Distribution directly to the parent Subscriber of the end-node target Subscriber.
4. If the Subscriber does not have a parent Subscriber, the Distributor checks its routing hierarchy for the Subscriber:
 - a. If the Subscriber is in the routing hierarchy, the Distributor uses that route to send the Distribution to the Subscriber.
 - b. If the Subscriber is not in the routing hierarchy, the Distributor sends the Distribution directly to the Subscriber.

In other words, if the Distributor can find a way to send the Distribution using its routing hierarchy, it uses the path in that hierarchy to get the Distribution to the Subscriber. Otherwise, it sends the Distribution directly to the Subscriber (or its parent Subscriber).

For that reason, you should make sure every Subscriber that regularly receives Distributions from a Distributor has some connection to the Distributor’s routing hierarchy. You can make this connection by listing a Subscriber in the hierarchy or by having one of the Subscribers in the hierarchy be its parent Subscriber.

You should generally not allow the Distributor to send Distributions over WAN links, except to such Subscribers that might be in the first tier of its routing hierarchy.

Consider the following in designing your Distributor’s routing hierarchy:

- ◆ **End-node Subscribers:** The only Subscribers that you need to add to the routing hierarchy are those you want to be used to pass on Distributions. End-node Subscribers that only receive Distributions and not pass them on do not need to be added to the routing hierarchy.

- ◆ **Configuring distribution routes:** To create the distribution routes, consider your network design and the number of Subscribers on each LAN. Then design the routing hierarchy to mimic your network topology.
- ◆ **Selecting multiple Subscribers:** During hierarchy creation, you can place multiple Subscribers at the same tier under a single Distributor or Subscriber.

IMPORTANT: The most efficient routing hierarchy is to have more tiers and fewer Subscribers per tier, than just a few tiers with many Subscribers per tier. Therefore, select only a few Subscriber servers per tier. This minimizes the workload for the Distributor or Subscriber server that is sending Distributions to other Subscriber servers. Tiering helps to share the workload of sending Distributions throughout the network.

- ◆ **Using multiple Distributors:** Multiple Distributors can use the same routing hierarchy of Subscribers, so that the same distribution route can be used by each Distributor.
- ◆ **Reusing Subscribers:** You should consider whether you might overload a Subscriber server if it should be a parent Subscriber in a routing hierarchy that services multiple Distributors.

Selecting Subscribers for the Distribution Routes

The purpose of the Distributor's routing hierarchy is to create the most efficient method for distributing to Subscribers. You need to determine which servers are best suited to be Subscribers in a routing hierarchy, and how many servers to include in the hierarchy.

Select a server that is robust in its physical configuration. For example, a fast CPU, plenty of RAM, and plenty of free hard disk space (especially on volumes other than sys: on NetWare servers).

Use the following criteria to determine which Subscribers to include in a Distributor's routing hierarchy:

- ◆ Is the Subscriber needed to minimize the Distributor's workload?
- ◆ Do you need other Subscribers to share the workload of a parent Subscriber on a given LAN?
- ◆ Is the Subscriber needed to minimize network traffic (such as through WANs or firewalls)?

To identify the Subscriber servers to use in a Distributor's routing hierarchy, create a list of the servers in your network that you want to use as parent Subscribers in a Distributor's routing hierarchy.

To help minimize network traffic, select at least one server on each LAN.

Identify the server objects that you want to be parent Subscribers in the Distributors' routing hierarchies:

CONFIGURATION PLANNING WORKSHEET

Under **item 16**, provide the names (including full context) for your parent Subscriber servers.

Configuring the Distribution Routes

Specify the following information on your network diagram:

CONFIGURATION PLANNING DIAGRAM

Write "parent=1" next to every location on the diagram that is separated from the Distributor's location by a WAN link or firewall (unless there is only one Subscriber at that location).

For every location on the diagram that requires additional parent Subscribers because of the high number of Subscribers, change "parent=1" to "parent=#" where # is the number of parent Subscribers the site needs for load-balancing.

Also note whether you want to use one parent Subscriber in a given location as the primary parent Subscriber (the only one at that location in the Distributor's routing hierarchy) for receiving Distributions and passing them on to other parent Subscribers in that location.

Be sure to include parent Subscribers at the Distributor's location, if needed.

Using the information from your network diagram, design your Distributors' routing hierarchies using the Subscribers you have selected:

CONFIGURATION PLANNING WORKSHEET

Under [item 15](#), create a hierarchy for each Distributor's routing hierarchy. You can reuse Subscriber servers in different Distributor's hierarchies.

1.1.7 Understanding Distribution Security

Server Management provides adequate security for Distributions that are sent within a secured network using certificates. However, Distributions could require additional security measures that are available in Server Management.

For more information about security, see [Chapter 7, "Security in Policy and Distribution Services," on page 299](#).

Review the following to determine whether you need any additional security for your Distributions:

- ◆ ["Determining Whether You Need Inter-Server Communications Security" on page 44](#)
- ◆ ["Determining Whether You Need Encryption Security for Windows Servers" on page 45](#)

Determining Whether You Need Inter-Server Communications Security

Policy and Distribution Services uses XMLRPC (Extensible Markup Language Remote Procedure Call) for its normal inter-server communications. XMLRPC optionally provides security for communicating securely across non-secured connections.

Policy and Distribution Services can use this security for inter-server communications between servers across non-secured connections, or between a management workstation and servers across non-secured connections. For example, firewalls, intranets, NAT configurations, and so on.

This inter-server communications security ensures that data received across a non-secured connection is from a trusted source, that it has not been tampered with en route, and that the data received can be trusted by other machines. This is accomplished through the use of signed security certificates and digital signatures.

This security requires modifications to certain text files, and is installed using a Server Management wizard.

The following are instances when you could want inter-server communication security:

- ♦ **ConsoleOne administration:** When you use a workstation to manage a Distributor server across a non-secured connection.
- ♦ **SET parameters:** When you create a SET Parameter policy or a software package for SET parameters, inter-server communication takes place to provide the target server's SET parameter information. This communication could cross a non-secured connection.
- ♦ **Server Down policy:** When you use this policy to down a server, the communication between the downed server and another server watching for it to come back up could cross a non-secured connection.

For more information, see [Section 7.3, "Security for Inter-Server Communication Across Non-Secured Connections,"](#) on page 313.

CONFIGURATION PLANNING WORKSHEET

Under [item 13](#), provide the NetWare and Windows server names where you need to install the inter-server communications security software.

Determining Whether You Need Encryption Security for Windows Servers

You normally do not need to encrypt Distributions that are sent within your secured network. However, you can use encryption to provide security for when you send Distributions outside your network. The NICI software is used for encrypting Distributions.

For some NetWare servers, NICI 2.6 is automatically installed with the operating system. However, version 2.6.4 is supported in ZENworks 7 Server Management. You may need to upgrade your NetWare version of NICI. Version 2.6.4 is shipped with ZENworks 7, and is also shipped with ZENworks for Servers 3.0.2 (including version 3 SP2).

For Windows, Linux, and Solaris servers, you must install NICI 2.6.4 on the Distributor and Subscriber servers where you expect encrypted Distributions to be built and extracted.

IMPORTANT: If you have NICI 2.4.6 running on your network, it is optional whether you upgrade to NICI 2.6.4, because these versions are compatible with each other.

If you need to install the NICI software on a Windows, Linux, and Solaris server, you must also install that same version on all Distributor and Subscriber servers in your network. Encryption does not work correctly if there are two different versions of NICI installed in your network.

For information on Distribution encryption, see [Section 7.2, "Distribution Security Using Encryption,"](#) on page 309.

CONFIGURATION PLANNING WORKSHEET

Under [item 14](#), provide the Windows, Linux, and Solaris server names where you need to install the NICI software.

1.1.8 Determining the Channels for the Distributions

Channels are used to group Distributions, to establish a schedule for passing a Distributor's Distributions on to Subscribers, and to list the Subscribers that are subscribed to the Channel so that the Distributor knows where to physically send the Distribution files.

You can create a Channel for a specific Distribution usage (such as virus pattern files, operating system support packs, or policy packages), or for a specific Distribution time (such as off-peak Distributions).

You can associate a Channel with Distributions from many Distributors. A Channel can be subscribed to by many Subscribers.

Subscribers subscribe to Channels in order to receive certain Distributions. Distributors associate their Distributions with the Channels so that the subscribed Subscribers can receive those Distributions.

If you are installing multiple Distributors, they can share Channels for their Distributions. For example, if Distributor A and Distributor B both want to send some of their Distributions to the same set of Subscribers, one Channel can be used by both Distributors.

Channels are used in providing Distributions to Subscribers. Consider the following:

- ◆ A Channel is not owned by any particular Distributor
- ◆ Distributors associate their Distributions with the Channels
- ◆ A Channel can have Distributions from multiple Distributors
- ◆ A Channel can be used to group related Distributions
- ◆ A Channel's schedule determines when the listed Distributions are sent
- ◆ A Subscriber subscribes to one or more Channels to receive all of the Distributions listed in those Channels
- ◆ A Subscriber cannot select an individual Distribution from the several that could be listed in a Channel (it must receive all of the Channel's Distributions)

In naming Channels, use a descriptive method. For example:

```
VirusProtect  
VProtectPatterns  
VirusProtection  
NW51patch4  
NW6patch1  
AUTOEXECNCF000326
```

You can manage your Channels more easily by:

- ◆ Using names that are purpose oriented
- ◆ Using a similar name for the Channel and its Distributions

CONFIGURATION PLANNING WORKSHEET

Under **item 21**, provide your Channel names. Make the names unique to help identify which Distributions they will hold.

You generally create a Channel for one or more related Distributions. However, for distribution flexibility, you can create one Channel for each application to be distributed.

CONFIGURATION PLANNING WORKSHEET

Under [item 22](#), provide the Distributions that belong to each Channel.

For ease of management, plan to create the Channel objects in the same context as your other Tiered Electronic Distribution objects, especially the Distribution objects.

CONFIGURATION PLANNING WORKSHEET

Under [item 20](#), provide the eDirectory context where the Channel object should be created.

1.1.9 Determining Subscribers' Subscriptions

You need to subscribe your Subscribers to Channels before they can receive their Distributions. This is done by subscribing a Subscriber or Subscriber Group to the Channel that is associated with the Distribution it needs:

- ◆ [“Subscribers” on page 47](#)
- ◆ [“Subscriber Groups” on page 47](#)

Subscribers

Because Subscribers do not access eDirectory, all configuration information in the Subscriber object's properties is pushed down to it from the configuring Distributor, if it is needed. This includes such information as working directory, log file level and location, console messaging level, variables, and so on.

Changes to a Subscriber object's properties are not in effect until the Distributor reads eDirectory again and sends a new Distribution with the configuration information down to the Subscriber.

For each Distribution, determine which Subscriber servers need a particular Distribution.

CONFIGURATION PLANNING WORKSHEET

Under [item 24](#), provide the Channel name for a Distribution (see [item 22](#)) and list the Subscribers that need that Distribution. Repeat for each Channel you provided in [item 21](#).

Subscriber Groups

A Subscriber Group is used for grouping Subscribers that have the same Distribution needs.

Subscriber Groups are useful when you are sending several different Distributions to the same set of Subscribers. There is no need to create a Subscriber Group if it is only associated with one Channel.

For example, Distribution A is in Channel A, Distribution B is in Channel B, and so on. Then, if you are not using a Subscriber Group, you need to subscribe each of your Subscribers to Channel A, then each to Channel B, and so on, which could be a very long process. However, by using a Subscriber Group, you only need to create the group, add the Subscribers to it, then subscribe that one group to each Channel.

Another use of a Subscriber Group is that when the group is associated with two or more Channels, you can edit the group's membership more easily than making the same changes in multiple Channels. For example, to remove a Subscriber from one Subscriber Group, you just edit that one group's properties. To remove that same Subscriber from several Channels, you need to edit each Channel's properties.

CONFIGURATION PLANNING WORKSHEET

Under [item 17](#), provide a unique name for the Subscriber Group.

Under [item 18](#), provide a list of Subscribers that need the same Distributions from the Channel (see [item 21](#) and [item 22](#)) where the group is subscribed.

Under [item 24](#), provide the Channel names for the Distributions that you want all of the Subscribers in the group to receive.

1.1.10 Determining the Distribution Schedules

Tiered Electronic Distribution has different schedules so that you can coordinate the various distribution processes. For more detailed information, see [Chapter 8, "Scheduling," on page 317](#).

Review the following to plan your Tiered Electronic Distribution schedules:

- ◆ ["Understanding Scheduling in Tiered Electronic Distribution" on page 48](#)
- ◆ ["Determining the Distributor's Refresh Schedule" on page 49](#)
- ◆ ["Determining the Distribution's Build Schedule" on page 49](#)
- ◆ ["Determining the Channel's Send Schedule" on page 49](#)
- ◆ ["Determining the Subscriber's Extract Schedule" on page 49](#)

Understanding Scheduling in Tiered Electronic Distribution

Both Tiered Electronic Distribution objects and individual Server Policies can be scheduled.

Tiered Electronic Distribution uses schedules to control when Distributors are refreshed and Distributions are built, sent, and extracted. Schedules do not affect the total resources used by a Distribution, but rather *when* the resources are used.

Some policies must be scheduled before they can be enforced. If you enable a policy, but do not schedule it, it is activated according to the schedule currently specified in the Default Package Schedule, which provides a default for scheduled policies. The default schedule is Run At System Startup.

If you configure several policies with the same schedule, the order they are run depends on the time stamps created when you created the policies. Therefore, when you view a list of policies, the order they are listed is the order that they are run.

If you want to control the order that certain policies are run, you should stagger their schedules, rather than rely on the time stamps to determine when they run. Therefore, consider the Tiered Electronic Distribution schedules you select when scheduling your policies, so that you do not have undesirable overlap, or out-of-sequence events that could cause some scheduled items to fail.

Other issues you might need to understand:

- ◆ How time zones can affect scheduling
- ◆ How policy schedules are affected by distribution schedules
- ◆ How distribution schedules can be affected by Distributor and Subscriber servers' non-Server Management software usage
- ◆ How the Randomly Dispatch option can affect scheduling
- ◆ How the Active and Inactive object options for the Tiered Electronic Distribution objects can affect scheduling and distribution flow

Determining the Distributor's Refresh Schedule

The Distributor's Refresh schedule determines when the Distributor should read eDirectory for new Distribution and Channel objects, or for configuration changes to existing Distribution and Channel objects. Upon a Distributor refresh, when the Build schedule starts the Distributor rebuilds the Distributions that it discovers to be new or changed, then sends them when the Send schedule starts.

The Refresh schedule is set to Never by default, which is recommended because an infinite loop could be encountered if the Refresh frequency is shorter than the time it takes to complete the building or sending of a Distribution. Therefore, you should normally refresh a Distributor manually.

If you want to use a different schedule than Never for Refresh, be certain that when the Distributor is refreshed it is not going to be in the middle of building or sending a Distribution.

As an example of when you might want to change the Refresh schedule from Never, if you create or change your Distributions daily and do not need to build and send them immediately, you can set the Refresh schedule to 1:00 AM daily to have your new Distribution objects or changes found by the Distributor so that it can build and send them during off-peak hours according to the Build and Send schedules.

Determining the Distribution's Build Schedule

The Build schedule determines when a Distributor is requested to build the individual pieces that comprise the Distribution.

During configuration, you are instructed to set each Distribution's Build schedule to allow the Distribution to be sent immediately after building it.

Determining the Channel's Send Schedule

The Send schedule provides a window of time for when a Distributor can send its Distributions to the Subscribers.

During configuration, you set each Channel's Send schedule to an interval of every 5 minutes, meaning that the Distributor can send its Distributions at any of the 5-minute intervals when the Channel's schedule fires.

Determining the Subscriber's Extract Schedule

The Extract schedule determines when a Subscriber can extract its received Distribution.

Before a Subscriber can use a Distribution that is sent to it, it must first extract the Distribution. Therefore, you should set the Subscriber's Extract schedule before you send the Distributions.

Determine when you want the various Subscriber servers to be active extracting Distributions. Depending on a Distribution's size, it could be best to have Distributions extracted during off-peak hours. For information on scheduling issues involving time zones, see [Section 8.2.5, "Scheduling Issues,"](#) on page 331, especially ["Calculating Time Differences"](#) on page 334.

CONFIGURATION PLANNING WORKSHEET

Under [item 23](#), provide the Subscribers' extract schedules.

1.2 Configuring Your Distribution System

Use these sections in the following order:

1. ["Installing Additional Distributors, Databases, and Subscribers"](#) on page 50
2. ["Setting Up Additional Distribution Security"](#) on page 54
3. ["Configuring the Distribution Flow"](#) on page 55
4. ["Creating the Distributions and Related Channels"](#) on page 57
5. ["Subscribing to the Distributions"](#) on page 59
6. ["Sending the Distributions"](#) on page 60

1.2.1 Installing Additional Distributors, Databases, and Subscribers

When installing Policy and Distribution Services for the first time, you installed one Distributor with a database file. If you planned to install more Distributors or databases (see ["Understanding Distributors"](#) on page 95 and [Section 10.2, "Determining How Many Databases You Need,"](#) on page 353), you should perform this installation now.

When installing Policy and Distribution Services for the first time, you might not have installed the Subscriber software to all of your servers. If you want to install the Subscriber software to more servers at this time, you should perform this installation now.

IMPORTANT: Any servers where you do not have the Subscriber software installed are not eligible to receive the Distributions you have planned to create and distribute at this time. However, when you install the Subscriber software to servers at a later date, you can subscribe them to existing Channels for receiving their Distributions.

To install additional Distributors, databases, and Subscriber software to more servers, do the following in order:

1. ["Preparing to Install"](#) on page 51
2. ["Starting the Installation Program"](#) on page 51
3. ["Selecting and Configuring the Distributor and Subscriber Servers"](#) on page 51
4. ["Completing the Installation"](#) on page 53

Preparing to Install

- 1 Make sure you have fulfilled all of the necessary requirements for your target Distributor and Subscriber servers.

For more information, see “[Server Requirements](#)” in the *Novell ZENworks 7 Server Management Installation Guide*.

- 2 Select the workstation you will use to install the Distributors and Subscribers.
- 3 If you have not already done so, log in to the eDirectory tree where you want to create the Server Management objects (worksheet [item 1](#)).

This should be the same tree where you extended the schema for ZENworks 7 Server Management.

You are automatically authenticated to all of the NetWare target servers in this tree during installation. You can select those servers, as well as servers in other trees or domains, for installing the Policy and Distribution Services software. However, this is the tree where all of the Server Management objects are installed for each of the selected servers.

- 4 Continue with “[Starting the Installation Program](#)” on page 51.

Starting the Installation Program

- 1 On the installation workstation, insert the *ZENworks 7 Server Management with Support Pack 1 Program CD*.

The startup screen is displayed. If the startup screen is not automatically displayed after inserting the CD, you can start it by running `winsetup.exe` at the root of the CD.

IMPORTANT: Installation from a CD in a remote server is not supported unless there is a drive mapped on the workstation to that remote server. For example, if you place the CD in a Windows server CD drive, then run the installation from a workstation, you must have a drive mapped on the workstation to the CD drive of that Windows server.

- 2 Select *Server Management*, then select *Policy-Enabled Server Management*.
This begins the installation program.
- 3 If you agree with the Software License Agreement, click *Accept > Next*.
- 4 On the Installation Type page, click *New Installation*, then click *Next*.
- 5 On the Installation Options page, make sure all three check boxes are selected.
- 6 On the eDirectory Tree for Creating Objects page, select the tree (worksheet [item 1](#)).
This is the tree where you initially created Server Management objects.
- 7 Continue with “[Selecting and Configuring the Distributor and Subscriber Servers](#)” on page 51.

Selecting and Configuring the Distributor and Subscriber Servers

- 1 On the Server Selection page, click *Add Server*.
- 2 Browse for and select the Distributor (worksheet [item 2](#)) and Subscriber (worksheet [item 3](#)) servers and click *OK*.
- 3 Configure each server listed on this page, then click *Next* to continue with the File Locations and Options page:

TIP: To quickly configure a specific role or set of roles for one or more servers, select the servers, right-click the selection, then select the role for the server. The options that apply to that role are automatically selected. Repeat for additional roles.

ZENworks Policy-Enabled Management Services

The following three options are all selected by default. If you want to install the Inventory Agent, you must also select to install the Policy and Distribution Server.

- ♦ **Policy and Distribution Services Server:** For each server that you want to be a Subscriber, select this check box.

For Tiered Electronic Distribution purposes, you can deselect the following:

Inventory Agents
Remote Management

Additional Options

The installation program detects whether these options are already installed on a target server and dims the option label. You can still select the check box to reinstall the component.

- ♦ **Distributor:** The Subscriber service is installed automatically to all target servers. Select this check box if you planned to make a Distributor server.
- ♦ **Server Management database:** This is the Policy and Distribution Services database that the Distributor logs to server (worksheet [item 4](#)). You should install it on the same server as the Distributor in order to minimize network traffic for database logging.

IMPORTANT: You can install the database to multiple servers per run of the installation program; however, you can only install one database per server. On the Database Settings page, you will be able to individually configure each database that is being installed. On the Database Logging page, you will identify which of the databases being installed is to be the one database for initial logging.

For Tiered Electronic Distribution purposes, you can deselect the following:

Inventory Database
Inventory Server
Inventory Proxy Server
ConsoleOne Snap-Ins

TIP: You can configure a group of selected servers with the same options by selecting the group and right-clicking the group. This displays the Custom Selection dialog box.

- 4 On the File Locations and Options page, do the following:
 - 4a For each Distributor server, edit the installation path if you do not want to use the default (worksheet [item 5](#)).

If you want all Distributor servers to have the same installation path, select all of the servers, then edit the path.
 - 4b For each Subscriber server, edit the installation path if you do not want to use the default (worksheet [item 6](#)).

If you want all Subscriber servers to have the same installation path, select all of the servers, then edit the path.

- 4c** To launch Policy and Distribution Services components on server startup, select the check box.
- 4d** To start services when the installation is finished, select the check box, then click Next.
- 5** On the Distributor Object Properties page, edit the properties as necessary (worksheet [item 7](#)), then click Next.
- 6** On the Subscriber Object Properties page, edit the properties as necessary (worksheet [item 8](#)), then click Next.
- 7** On the Database Settings page, do the following:
 - 7a** Edit the database file's path if you do not want to use the default (worksheet [item 9](#)).
Because the database file can become very large, we recommend that you change the default NetWare volume from sys: to another volume on that server.
 - 7b** Edit the Database object's name, if desired (worksheet [item 10](#)).
 - 7c** Change the Database object's container, if desired (worksheet [item 11](#)).
- 8** If you chose to install the Policy and Distribution Services database, the Log to a Server Management Database That Will Be Installed option is selected; click Next to display the Summary page.
- 9** Continue with ["Completing the Installation" on page 53](#).

Completing the Installation

- 1** To save the current installation configuration for future use in installing Distributors, on the Summary page select the *Save the following* check box.
- 2** Provide a path and filename for the template file.
If you attempt to quit the installation program without clicking Finish, you are prompted to save your current installation configuration to an installation template file.
You can reuse this template to speed up filling in installation pages in subsequent installations of Distributors or Subscribers.
- 3** Click *Finish* to begin the installation process.
- 4** After the installation program has finished, review the installation log file to determine whether any components failed to install.
The log file is located at:
`%TEMP%_resnumber.txt`
where *number* is a three-digit number that is increased incrementally each time a new installation log is created.
- 5** If necessary, rerun the installation program.
Select only the components that failed to install.
- 6** Rerun the installation program once for each additional database that needs to be installed (worksheet [item 4](#)).
On the Server Selection page add only one of the Distributors where you planned to have a database installed, but have not installed it yet. Then, click only the Database column for that database's Distributor server and fill in the applicable information on the remaining installation pages.

- 7 To set up additional distribution security, continue with [Section 1.2.2, “Setting Up Additional Distribution Security,” on page 54](#); otherwise, continue with [Section 1.2.3, “Configuring the Distribution Flow,” on page 55](#).

1.2.2 Setting Up Additional Distribution Security

To ensure that you have the proper security for your Distributions, do the following tasks that are applicable:

- ♦ [“Installing NICI 2.6.4” on page 54](#)
- ♦ [“Setting Up Inter-Server Communications Security” on page 55](#)

Installing NICI 2.6.4

If you need Distribution encryption support for certain NetWare, Windows, Linux, or Solaris Subscriber servers, NICI 2.6.4 is supported in ZENworks 7 Server Management. If not, skip to [“Configuring the Distribution Flow” on page 55](#).

If you previously updated your servers to NICI 2.6.4 using ZENworks for Servers 3 SP2, skip to [“Configuring the Distribution Flow” on page 55](#).

IMPORTANT: All servers that are sending or receiving encrypted Distributions must be running the same version of NICI. Otherwise, encrypted Distributions to any of those servers will fail.

You must install NICI 2.6.4 to all Subscribers subscribed to the Channel that you select for the software package used to distribute NICI. NICI 2.6.4 must also be running on any Distributor server that creates encrypted Distributions.

However, if you already have NICI 2.4.6 installed, it is optional whether you upgrade to NICI 2.6.4, because these versions are compatible with each other.

A NICI update is contained on the *ZENworks 7 with Support Pack 1 Companion 2* CD, which is installed to Windows servers using the Novell International Cryptographic Infrastructure (NICI) menu option.

A software package update for NICI 2.6.4 is also provided on the *ZENworks 7 with Support Pack 1 Companion 2* CD.

When you install NICI 2.6.4, the installation program does not check to see if NICI is already installed.

Select the appropriate installation method:

- ♦ [“Installing NICI on Windows Servers” on page 54](#)
- ♦ [“Installing NICI Using the Server Software Package” on page 55](#)

Installing NICI on Windows Servers

To install NICI 2.6.4 on Windows servers:

- 1 On a Windows workstation, insert the *ZENworks 7 with Support Pack 1 Companion 2* CD.
- 2 Select the *Companion Programs and Files* option, then click *more >>* to access the *Companion 2* CD menu.

- 3 Select the *Novell International Cryptographic Infrastructure (NICI)* menu option.
- 4 Follow the installation instructions.
- 5 Continue with [Section 1.2.3, “Configuring the Distribution Flow,” on page 55.](#)

Installing NICI Using the Server Software Package

To install NICI 2.6.4 on any supported server:

- 1 On a Windows workstation, insert the *ZENworks 7 with Support Pack 1 Companion 2 CD*.
- 2 Copy the `nic_i265.exe` file from the `\NICI` directory on the CD to a location on your workstation, then extract the file.
- 3 Copy the `nic_i264.cpk` file that was extracted to a location on the Distributor server where you create the Software Package Distribution for installing NICI 2.6.4.
- 4 Create and send the Distribution to each Subscriber server where encrypted Distributions are received.

For information on creating and sending Software Package Distributions, see [Section 3.4.4, “Creating a Distribution,” on page 123.](#)

- 5 Continue with [Section 1.2.3, “Configuring the Distribution Flow,” on page 55.](#)

Setting Up Inter-Server Communications Security

If you are distributing to servers outside your secured network (worksheet [item 13](#)), see [Section 7.3, “Security for Inter-Server Communication Across Non-Secured Connections,” on page 313](#) for detailed instructions on setting up security for inter-server communications.

1.2.3 Configuring the Distribution Flow

You need to configure your distribution system to ensure the most efficient use of your network in sending Distributions by setting up the Distributors’ routing hierarchies. This was not done for any Distributor when you installed Policy and Distribution Services.

To configure your distribution system:

- ♦ [“Configuring the Distributor Routing Hierarchies” on page 55](#)
- ♦ [“Configuring Parent Subscribers” on page 56](#)
- ♦ [“Configuring Subscriber Groups” on page 56](#)

Configuring the Distributor Routing Hierarchies

- 1 In ConsoleOne, right-click a Distributor object (worksheet [item 2](#)), then click *Properties*.
- 2 Select the *Routing* tab and do the following:
 - 2a Click *Add* and browse for your first tier Subscriber servers (worksheet [item 15](#)), then click *Select > OK*.

This sets up your first tier of Subscriber servers. These receive Distributions directly from the Distributor.
 - 2b Select one of the Subscriber servers in the first tier of the routing tree, click *Add* and browse for your next tier of Subscriber servers to go under that first tier Subscriber (worksheet [item 15](#)), then click *Select > OK*.

This sets up a second tier of Subscriber servers for the one Subscriber that you selected. These second-tier Subscribers receive Distributions indirectly from the Distributor via the Subscriber server above them in the hierarchy.

- 2c** Repeat **Step 2b** for each of the first-tier Subscribers until you have selected all of the second-tier Subscribers for this part of the hierarchy.
 - 2d** Select one of the Subscriber servers in the second tier of the routing tree, click *Add* and browse for your next tier of Subscriber servers to go under that Subscriber (worksheet **item 15**), then click *Select > OK*.
 - 2e** Repeat **Step 2d** for each of the second tier Subscribers until you have selected all of the third-tier Subscribers for this part of the hierarchy.
 - 2f** Repeat this process, tier by tier, until you have completed your planned routing hierarchy for the current Distributor.
- 3** Repeat **Step 1** through **Step 2** for your other Distributors.
 - 4** When you have finished building the routing hierarchy, click *OK*.
 - 5** Continue with “**Configuring Parent Subscribers**” on page 56.

Configuring Parent Subscribers

All Subscribers should not receive their Distributions directly from a Distributor. The Distributor’s routing hierarchy provides a way to minimize the Distributor’s workload in sending Distributions.

For Subscriber servers to receive their Distributions using the routing hierarchy, you need to identify a parent Subscriber that is in the routing hierarchy for each end-node Subscriber (the Subscriber to receive the Distribution). This allows an end-node Subscriber to receive its Distributions through the routing hierarchy, rather than directly from a Distributor.

A Subscriber that is in the Distributor’s routing hierarchy does not need to have a parent Subscriber in order to receive a Distribution from that Distributor. Distributors check their routing hierarchies first, then check for parent Subscribers second.

To associate Subscribers with parent Subscribers:

- 1** In ConsoleOne, select a group of Subscriber objects for servers that you planned to have serviced by a particular parent Subscriber (worksheet **item 16**), right-click the selected group, click *Properties of multiple objects*, in the *Parent Subscriber* field browse for the parent Subscriber object, then click *OK > OK*.

Because you can do multiple editing of eDirectory objects, you can select all of the Subscribers that are serviced by one parent Subscriber and edit the Parent Subscriber field once for all of them.

- 2** Repeat this process for all end-node Subscribers.
- 3** Continue with “**Configuring Subscriber Groups**” on page 56.

Configuring Subscriber Groups

To create and populate a Subscriber Group:

- 1** In ConsoleOne, select the container to hold the Subscriber Group object, click *File > New > Object*, then select *TED Subscriber Group*.

- 2 In the New TED Subscriber Group dialog box, specify a *Subscribe Group* name (worksheet [item 17](#)), click *Define additional properties*, then click *OK*.
- 3 Click *General > Settings* and provide a description.
- 4 To populate the group with Subscribers, select the *Members* tab, then do the following:
 - 4a Click *Add*, browse for and select the Subscriber objects (worksheet [item 18](#)), then click *OK*.
 - 4b To remove any Subscribers from the list, select the Subscribers and click *Delete*.
 - 4c To view the properties of any Subscriber, select the Subscriber and click *Details*.
- 5 Click *OK* when you have finished configuring the Subscriber Group object.
- 6 Continue with [Section 1.2.4, “Creating the Distributions and Related Channels,” on page 57](#).

1.2.4 Creating the Distributions and Related Channels

The following are generic instructions for creating a Distribution. For more detailed instructions for most Distribution types, see [Chapter 3, “Tiered Electronic Distribution,” on page 85](#). For steps on using the Distribution Wizard to create a File or FTP Distribution, see [Section 3.4.12, “Using the Distribution Wizard,” on page 143](#).

You first need to create the Distribution, then create the Channel (if you don’t use an existing Channel):

- ◆ [“Creating and Configuring the Distribution” on page 57](#)
- ◆ [“Creating and Configuring the Channel” on page 58](#)

Creating and Configuring the Distribution

- 1 In ConsoleOne, locate the containers where the Tiered Electronic Distribution objects were installed.
- 2 Right-click the container for Distributions, click *New > Object*, then select *TED Distribution*.
- 3 Specify a Distribution name (worksheet [item 19](#)).
Name the Distribution so you can identify what it contains.
- 4 Browse to and select the Distributor object to own this Distribution (worksheet [item 19](#)).
Each Distribution is associated with a single Distributor. That Distributor is responsible for building and sending the Distribution.
- 5 Select the *Define additional properties* check box.
- 6 Click *OK* to create the object.
The properties for the Distribution are now displayed.
- 7 Select the *Type* tab; in the Select Type drop-down box, select a Distribution type (worksheet [item 19](#)).
- 8 Configure the Distribution.
For information on configuring the different Distribution types, see [Section 3.4, “Distributions,” on page 110](#).
Use the up-arrow and down-arrow buttons to change the distribution order.
- 9 Select the *Schedule* tab.

The Distribution's schedule determines how often the Distributor attempts to build a new version of the Distribution. A new version is built only if there have been changes since the last version was built.

- 10 Select *Run Immediate* from the drop-down list.

This causes the Distributor to build the Distribution as soon as it reads eDirectory for the Distribution information.

- 11 Click *OK* at the bottom of the Distribution Properties dialog box to save all changes.

- 12 If you have not previously resolved certificates, for NetWare and Windows servers, select *Yes* when prompted to copy security certificates.

For Linux and Solaris servers, certificates must be resolved manually if you do not have a drive mapped to them. For more information, see [Section 7.1.6, "Resolving Certificates," on page 303](#).

The Distributor needs to have been run at least once so that its certificates can be minted (created).

A Distributor needs to resolve its certificates only once per Subscriber.

The Subscriber software does not need to be running on the server for security certificates to be resolved. The server only needs to be up.

ConsoleOne sends security certificates to each Subscriber server that subscribes to the Channel that was selected in the Channel Tab. Each Subscriber must have a security certificate from the Distributor before it can receive Distributions from that Distributor.

It can take several minutes to copy a security certificate to each Subscriber.

IMPORTANT: Certificate copying only needs to be done once for each Distributor/Subscriber relationship.

- 13 If you receive an error when the Distributor tries to copy to a Windows Subscriber, enter the following for the path:

```
\\IP_Address\zen$\pds\ted\security
```

where *IP_Address* is the IP address of that Windows Subscriber.

- 14 If you receive an error when the Distributor tries to copy to a Linux or Solaris Subscriber, or you cannot browse for the Server to select it for resolving certificates, you must map a drive to the server (such as through using Samba), and then repeat resolving certificates.

- 15 Repeat these steps for any other Distributions you want to create at this time (worksheet [item 19](#)).

- 16 Continue with ["Creating and Configuring the Channel" on page 58](#).

Creating and Configuring the Channel

Channel objects are used to associate Subscribers with Distributions. When Subscribers subscribe to a Channel, they receive all of the Distributions associated with that Channel. Each Channel has a schedule that determines when the Distributions associated with it are to be sent to the Subscribers.

- 1 In ConsoleOne, locate the container where the Channel objects reside (worksheet [item 20](#)).

This container should already exist.

We suggest for ease of management that you use the same OU for all Channels.

- 2 Right-click the Channel object's container, click *New > Object*, select *Channel*, then click *OK*.

- 3 Specify a name for the Channel (worksheet [item 21](#)) and click *OK*.
You could name your Channels according to the Distributions you intend for them. For example, Channel - Antivirus Update.
- 4 Right-click the new Channel object and click *Properties*.
- 5 Select the *Distributions* tab, click *Add*, browse for and select the Distributions for the Channel (worksheet [item 22](#)), then click *OK*.
This associates the Distributions with the Channel. The Subscribers that are subscribed to this Channel receive all of the Distributions currently listed there.
- 6 To set the Channel's Send schedule, select the *Schedule* tab, select *Interval*, specify the interval as every hour, then click *OK*.
- 7 Repeat [Step 1](#) through [Step 5](#) for each Channel you have planned (worksheet [item 21](#)).
- 8 Continue with [Section 1.2.5, "Subscribing to the Distributions,"](#) on page 59.

1.2.5 Subscribing to the Distributions

- ♦ ["Setting Subscribers' Extract Schedules"](#) on page 59
- ♦ ["Subscribing to the Channels"](#) on page 59

Setting Subscribers' Extract Schedules

Before a Subscriber can use a Distribution that is sent to it via Tiered Electronic Distribution, it must extract the Distribution. Therefore, the Subscriber's extraction schedule must be set before sending the Distributions.

- 1 In ConsoleOne, right-click the Subscriber object (worksheet [item 23](#)) for a server where you want to set the extraction schedule, then click *Properties*.
- 2 Select the *Schedule* tab, click the arrow for the drop-down box, select *Run Immediately*, then click *OK*.
This causes the selected Subscriber to extract its Distributions as soon as they are received.
- 3 Repeat [Step 1](#) and [Step 2](#) as necessary until all Subscriber schedules have been set.
- 4 Continue with ["Subscribing to the Channels"](#) on page 59.

Subscribing to the Channels

Subscribers must subscribe to a Channel in order to receive the Distributions associated with that Channel. In the following steps, you will associate all of your Subscribers to the Channels created previously.

- 1 In ConsoleOne, right-click a Channel object (worksheet [item 21](#)) and click *Properties*.
- 2 Select the *Subscribers* tab, click *Add*, browse for each of the Subscriber or Subscriber Group (worksheet [item 24](#)) objects to be subscribed to this Channel, click *Select*, then click *OK*.
- 3 Select the *General* tab and make sure the *Active* check box is selected.
- 4 Click *OK* to close the Channel object's properties and save the changes.
- 5 Select *No* when prompted to copy security certificates.
- 6 Repeat [Step 1](#) through [Step 5](#) for each Channel (worksheet [item 21](#)).
- 7 Continue with [Section 1.2.6, "Sending the Distributions,"](#) on page 60.

1.2.6 Sending the Distributions

Now that you have installed, created, and configured your Distributors, Subscribers, Channels, and Distributions, you can begin the Distribution process.

Do the following in order:

1. [“Scheduling the Distribution and Refreshing the Distributor” on page 60](#)
2. [“Verifying That the Distribution Process Was Successful” on page 60](#)

Scheduling the Distribution and Refreshing the Distributor

- 1 In ConsoleOne, right-click the Distributor object (worksheet [item 2](#)) and click *Properties*.
- 2 On the Distribution object’s *Build Schedule* tab, click *Send Distribution immediately after building*, then click *OK* to close the properties.

The Distribution is sent as soon as it is built, regardless of the Channel’s Send schedule.

- 3 Right-click the Distributor object and click *Refresh Distributor*.

This causes the Distributor to read eDirectory and obtain all of the changes that were made in eDirectory. The manual refresh of the Distributor is the recommended method. For more information, see [“Determining the Distributor’s Refresh Schedule” on page 49](#).

- 4 Continue with [“Verifying That the Distribution Process Was Successful” on page 60](#).

Building the Distribution begins immediately (according to the Build schedule you set previously). The Distribution is sent within five minutes (according to the Send schedule you set previously).

As soon as the Subscribers receive the entire Distribution, they extract the contents to the Subscriber’s working directory that you specified in the Subscriber object’s properties.

Verifying That the Distribution Process Was Successful

There are a number of ways you can verify that your Distribution process has worked:

- ♦ **iManager:** The Tiered Distribution View and Subscriber Distribution View are the easiest methods for determining this information. For help on using those views, access the iManager Help on those pages.
- ♦ **Reporting:** Run a report on the Distribution to see its status. For information on Tiered Electronic Distribution reporting, see [Chapter 11, “Reporting,” on page 363](#).
- ♦ **Log files:** Depending on the logging levels you are using, you can review the log files for distribution statuses. The log files (.log) can be found in the Distributors’ and Subscribers’ [working directories](#).
- ♦ **Distribution files:** Compare the Distribution file on the Distributor’s file system (under `\zenworks\pds\ted\dist`) with the Subscriber’s file system (under `\zenworks\pds\ted\sub\individual_Distribution's_path`) to see if it was received. The Distribution file uses the same name on both servers.

1.3 Managing Your Distribution System

Your Policy and Distribution Services system is now set up and ready for use. You can revisit [Section 1.2, “Configuring Your Distribution System,” on page 50](#) at any time and use the applicable sections to update your distribution system.

You can manage your distribution system using the ConsoleOne and iManager tools. There is some functionality in one tool that is not in the other. Generally, you can use ConsoleOne for installation and setup tasks, and iManager for management tasks. For more information, see [Section 2.5, “Comparing the ZENworks Server Management Role in iManager with ConsoleOne Capabilities,”](#) on page 82.

For information on using ConsoleOne, see the following:

- ◆ [Chapter 3, “Tiered Electronic Distribution,”](#) on page 85
- ◆ [Chapter 4, “Server Policies,”](#) on page 193
- ◆ [Chapter 5, “Server Software Packages,”](#) on page 237
- ◆ [Chapter 6, “Desktop Application Distribution,”](#) on page 273
- ◆ [Chapter 11, “Reporting,”](#) on page 363

For information on using iManager, see [Chapter 2, “Novell iManager,”](#) on page 63.

If you have not yet installed Novell® iManager, see “[Management-Specific Workstation Requirements](#)” in the *Novell ZENworks 7 Server Management Installation Guide*. ZENworks® 7 Server Management supports iManager 2.0.2 and 2.5/2.6. However, version 2.5 or later is required for Windows Server 2003.

The ZENworks Server Management role in iManager enables you to manage Server Policies and Tiered Electronic Distribution objects, agents, and processes from any location where the Web browser Internet Explorer 6 SP1 or later is available. The Server Management plug-ins to iManager only work in this browser. Other Web browsers are not supported in ZENworks 7.

Using the ZENworks Server Management role, you can:

- ◆ Create, modify, and delete Tiered Electronic Distribution objects (Distribution, Subscriber, Distributor, Channel, Subscriber Group, and External Subscriber).
- ◆ Create, modify, delete, distribute, and enforce policies and policy packages.
- ◆ View a graphical representation of your distribution system, which makes it easy to track a Distribution from Distributor to end-node Subscriber, no matter how many parent Subscribers the Distribution passes through.
- ◆ Display a browser-based console, called the Remote Web Console, for each Distributor Agent and Policy/Package Agent in your system. From the Remote Web Console, you can check the configuration of any agent, monitor the activities of any agent, and control many agent functions, such as forcing an action on a Distributor or Subscriber server to happen immediately, and monitoring the status of a Distribution or Subscriber.

The following sections help you make the most of the features available to you in the ZENworks Server Management role:

- ◆ [Section 2.1, “Accessing the ZENworks Server Management Role in iManager,” on page 63](#)
- ◆ [Section 2.2, “Managing Tiered Electronic Distribution Objects,” on page 67](#)
- ◆ [Section 2.3, “Monitoring the Distribution Process,” on page 69](#)
- ◆ [Section 2.4, “Managing the Agents through Remote Web Console,” on page 72](#)
- ◆ [Section 2.5, “Comparing the ZENworks Server Management Role in iManager with ConsoleOne Capabilities,” on page 82](#)

2.1 Accessing the ZENworks Server Management Role in iManager

Review the following sections to log in to iManager and to become familiar with ZENworks role in iManager:

- ◆ [Section 2.1.1, “Logging in to iManager,” on page 64](#)
- ◆ [Section 2.1.2, “Becoming Familiar with the Interface,” on page 65](#)
- ◆ [Section 2.1.3, “Viewing the Roles and Tasks,” on page 65](#)

2.1.1 Logging in to iManager

The steps to log in to iManager are different for versions 2.0.2 and 2.5/2.6, because version 2.0.2 is tree-dependent and version 2.5/2.6 is not.

To access iManager in your Web browser:

- ♦ “Logging in to iManager 2.0.2” on page 64
- ♦ “Logging in to iManager 2.5/2.6” on page 64

Logging in to iManager 2.0.2

- 1 If you are not logged in to the Novell eDirectory™ tree where Tiered Electronic Distribution objects are located, log in.

TIP: If you are running iManager on a Windows server where the Novell Client™ is not installed, specify the IP address of a server where a replica of your eDirectory tree resides, instead of providing the tree name itself.

- 2 Access the following URL:

`http://server/nps/iManager.html`

where *server* is the IP address or DNS hostname of the server where iManager is installed.

The following dialog is displayed:



- 3 If the iManager login page does not appear, make sure that you entered the correct server designation and that you entered `nps` and `iManager.html` exactly as shown in the example, because it is case sensitive.

TIP: You might need to use `https` instead of `http`.

- 4 Provide the username and password for the server that you identified in **Step 2**, then press Enter or click *Login*.

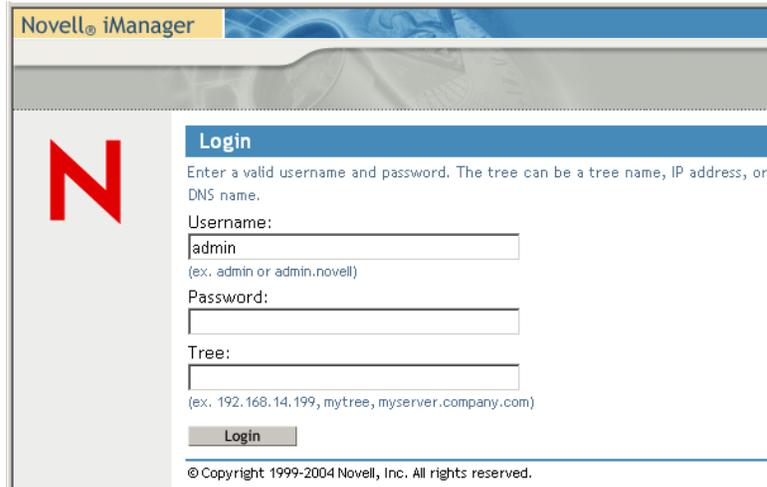
Logging in to iManager 2.5/2.6

- 1 Access the following URL:

`http://server/nps/iManager.html`

where *server* is the IP address or DNS hostname of the server where iManager is installed.

The following dialog is displayed:



The image shows the Novell iManager login dialog box. It has a title bar that says "Novell iManager". On the left side, there is a large red letter "N". The main area is titled "Login" and contains the following text: "Enter a valid username and password. The tree can be a tree name, IP address, or DNS name." Below this text are three input fields: "Username:" with the value "admin" and a hint "(ex. admin or admin.novell)", "Password:" (empty), and "Tree:" (empty) with a hint "(ex. 192.168.14.199, mytree, myserver.company.com)". At the bottom of the form is a "Login" button. Below the form, there is a copyright notice: "© Copyright 1999-2004 Novell, Inc. All rights reserved."

- 2 If the iManager login page does not appear, make sure you entered the correct server designation and that you entered `nps` and `iManager.html` exactly as shown in the example, because it is case sensitive.

TIP: You might need to use `https` instead of `http`.

- 3 Provide the username and password for the server that you identified in [Step 1](#).
- 4 Enter the tree designation for that server, then press Enter or click *Login*.

iManager 2.5/2.6 is not tree-dependent. Therefore, you can specify the tree during login, instead of before logging in (as in version 2.0.2), and can identify the tree using either the IP address of a server, a tree name, or the DNS name of a server.

2.1.2 Becoming Familiar with the Interface

- 1 After you successfully log in, the main iManager page is displayed. The top frame provides icons that represent its features:



- 2 Move the mouse pointer over the icons to review the purpose their functions. The mouse-over text appears to the right of the row of icons. By default, the *Roles and Tasks* icon is active, which is where the ZENworks functions reside.

2.1.3 Viewing the Roles and Tasks

- 1 **Open iManager**, then select the *Roles and Tasks* icon.

By default, the ZENworks Server Management role should be displayed in the left pane at the bottom of the tree structure.

2 In the left panel, expand *ZENworks Server Management* to list the available tasks:



These tasks provide the following functionalities:

| Task | Functionality |
|------------------------------|---|
| Create TED Object | Create any Tiered Electronic Distribution object, except a Distributor or Subscriber. |
| Delete TED Object | Delete any Tiered Electronic Distribution object. |
| Edit TED Object | Edit the properties of any Tiered Electronic Distribution object. |
| Remote Web Console | Viewing and managing Tiered Electronic Distribution information or Server Policies information. |
| Subscriber Distribution View | Viewing and managing selected Subscribers and all of their Distributions. |
| Tiered Distribution View | Viewing and managing selected Distributions or Distributors and all of their Distributions. |

3 If some of the above ZENworks Server Management tasks are not displayed, and you have Role-Based Services (RBS) configured, you might need to upgrade or reinstall the ZENworks Server Management module for the administrators who need access to the missing tasks.

For example, after upgrading the ZENworks Server Management plug-ins for iManager, if a new task was introduced by the upgrade, it will not be displayed for the RBS collections that are configured.

To solve this, follow the steps applicable to the version of iManager you are using:

- ♦ [“Upgrading Collections in iManager 2.0.2” on page 66](#)
- ♦ [“Reinstalling Collections in iManager 2.5/2.6” on page 67](#)

4 Continue with the task that you want to perform:

- ♦ [Section 2.2, “Managing Tiered Electronic Distribution Objects,” on page 67](#)
- ♦ [Section 2.3, “Monitoring the Distribution Process,” on page 69](#)
- ♦ [Section 2.4, “Managing the Agents through Remote Web Console,” on page 72](#)

Upgrading Collections in iManager 2.0.2

- 1 [Open iManager](#), then click the *Configure* icon.
- 2 Under *RBS Configuration*, click *Configure iManager*.
- 3 Select *Upgrade Collections*, then click *Next*.
- 4 Make sure the collections you want to upgrade are selected, then click *Next*.

Only the collections that have out-dated or previously not installed iManager modules are displayed.

- 5 Make sure that ZFSModule is selected, select the scope, then click *Start*.
- 6 Click *Close* after the module is shown to be successfully upgraded.
- 7 Click the *Roles and Tasks* icon.

The missing ZENworks Server Management roles should now be displayed under *ZENworks Server Management*.

- 8 Continue with [Step 4 on page 66](#).

Reinstalling Collections in iManager 2.5/2.6

- 1 [Open iManager](#), then click the *Configure* icon.
- 2 Expand *Role Based Services*, then click *RBS Configuration*.
- 3 Under the *Name* column, select the desired collection to edit.
- 4 Under the *Name* column, select the *ZENworks Server Management* role.
- 5 Under the *Reinstall* column, click the check box for the listed ZFSModule name.
- 6 Click *Reinstall* (the column heading).
- 7 Click *OK* in response to the information dialog box to reinstall the module.
- 8 After the module is shown to be successfully reinstalled, click the *Roles and Tasks* icon.

The missing ZENworks Server Management roles should now be displayed under *ZENworks Server Management*.

- 9 Continue with [Step 4 on page 66](#).

2.2 Managing Tiered Electronic Distribution Objects

Acting in the ZENworks Server Management role in iManager, you can create, edit, or delete some of the following Tiered Electronic Distribution objects in eDirectory:

Distributor (cannot create)
Channel
Distribution
Subscriber (cannot create)
Subscriber Group
External Subscriber

For these Tiered Electronic Distribution objects, you can perform all of the same management tasks in iManager that you can perform in ConsoleOne®:

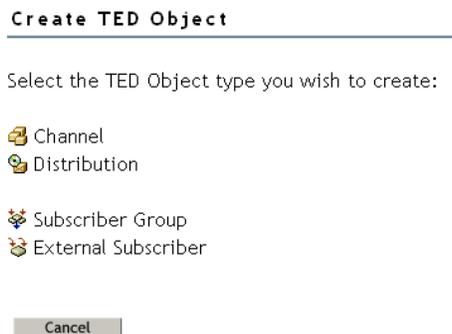
- ◆ [Section 2.2.1, “Creating Tiered Electronic Distribution Objects in iManager,” on page 68](#)
- ◆ [Section 2.2.2, “Editing Tiered Electronic Distribution Object Properties in iManager,” on page 68](#)
- ◆ [Section 2.2.3, “Deleting Tiered Electronic Distribution Objects in iManager,” on page 69](#)

The following Policy and Distribution Services management tasks cannot be performed in iManager and must be performed using ConsoleOne:

- ♦ Managing the Server Management database. See [Chapter 10, “ZENworks Database,” on page 351](#)
- ♦ Generating reports from the Server Management database. See [Chapter 11, “Reporting,” on page 363](#)

2.2.1 Creating Tiered Electronic Distribution Objects in iManager

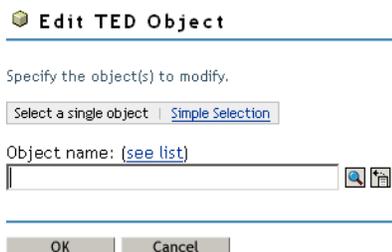
- 1 **Open iManager**, select *Roles and Tasks* in the top frame, expand *ZENworks Server Management* in the left frame, then click *Create TED object*.



- 2 Select the type of object you want to create.
Any Distribution type you can create in ConsoleOne, you can also create in iManager.
- 3 Provide the information required for that object type, such as a unique name for the object, the context where you want to create the object, and so on.
Click the *Help* icon (question mark) for more information.
- 4 Click *OK* to finish creating the object.
- 5 Continue with [Section 2.2.2, “Editing Tiered Electronic Distribution Object Properties in iManager,” on page 68](#) to configure the new Tiered Electronic Distribution object.

2.2.2 Editing Tiered Electronic Distribution Object Properties in iManager

- 1 **Open iManager**, select *Roles and Tasks* in the top frame, expand *ZENworks Server Management* in the left frame, then click *Edit TED Object*.



- 2 Browse to and select the Tiered Electronic Distribution object whose properties you want to edit, then click *OK*.

The same property pages and options are available in iManager that are available in ConsoleOne.

You can click *Help* on each property page for information on setting the options.

- 3 Configure the object as needed, then click *OK* to save the new properties settings.

2.2.3 Deleting Tiered Electronic Distribution Objects in iManager

- 1 Open **iManager**, select *Roles and Tasks* in the top frame, expand *ZENworks Server Management* in the left frame, then click *Delete TED Object*.



- 2 Browse to and select one or more Tiered Electronic Distribution objects to delete, then click *OK* to list the objects on the Delete TED Objects page.
- 3 Click the *Help* icon for information about the repercussions of deleting specific types of objects from your distribution system.
- 4 Click *OK* to delete the listed objects, then click *OK* again to confirm.
- 5 Follow any instructions in the online help to reconfigure remaining objects so that the deletion does not disrupt your distribution system.

2.3 Monitoring the Distribution Process

The Tiered Distribution View enables you to track a Distribution from its Distributor through any parent Subscribers down to the end-node Subscriber. This helps you determine which Subscribers have received the Distribution, where they received it from, and when they received it. This, in turn, helps you troubleshoot and correct any problems that might occur during the distribution process.

The Subscriber Distribution View provides a status view of all Distributions for each Subscriber that you add to a watch list. You can use this view to troubleshoot a Subscriber's Distributions.

These capabilities are not available in ConsoleOne.

The following sections explain how to use these views:

- ♦ [Section 2.3.1, "Using the Tiered Distribution View," on page 69](#)
- ♦ [Section 2.3.2, "Using the Subscriber Distribution View," on page 70](#)

2.3.1 Using the Tiered Distribution View

To access the Tiered Distribution View in iManager:

- 1 **Open iManager**, select *Roles and Tasks* in the top frame, expand *ZENworks Server Management* in the left frame, then click *Tiered Distribution View*.
- 2 Browse to and select the Distribution you want to track, then click *Next*.
- 3 Select the Channel through which you want to track the Distribution, then click *Next*.
The Distribution System window lists Subscribers that should receive the Distribution.
- 4 Click *Expand All* to display the routing hierarchy between the Distributor that built and sent the Distribution and the end-node Subscribers that should have received it.

or

Select an individual server to expand its part of the hierarchy.

Icons indicate the status of the Distribution:

| Icon | Meaning |
|---|---|
|  | The Distribution has been received and extracted successfully. |
|  | The Distribution has been received but not yet extracted. Check the Subscriber's extract schedule to see whether extraction has been attempted. If extraction was attempted and failed, check the Subscriber's event log to see what error occurred during extraction. See Section 2.4.2, "Managing the Distributor Agent," on page 77 . |
|  | The Distribution was not successfully received by the Subscriber. Check the Subscriber's event log for an error message describing the problem. See Section 2.4.2, "Managing the Distributor Agent," on page 77 . |
|  | The Distributor has not received any response from the Subscriber concerning the status of the Distribution. Check the status of the Subscriber and any parent Subscribers between it and the Distributor. See Section 2.4.2, "Managing the Distributor Agent," on page 77 . |

- 5 Fill in the time space in the *Refresh screen every ___ seconds* field, then click *Start* to refresh the display at that frequency.
Only seconds can be entered.
This is useful for troubleshooting the distribution process as it occurs.
- 6 To display status information, select a Distributor or Subscriber, then click *Remote Web Console*.
For information about the types of status information you can obtain, see [Section 2.4, "Managing the Agents through Remote Web Console," on page 72](#).
- 7 To check configuration information, select a Distributor or Subscriber, then click *eDirectory Configuration*.
You can edit the Distributor or Subscriber object properties just as if you had clicked *Edit TED Object* under *ZENworks Server Management*. The same property pages and options are available in iManager that are available in ConsoleOne.

2.3.2 Using the Subscriber Distribution View

To access the Subscriber Distribution View in iManager:

1 **Open iManager**, select *Roles and Tasks* in the top frame, expand *ZENworks Server Management* in the left frame, then click *Subscriber Distribution View*.

2 Click *Add* to select the Subscribers that you want to track.

Each Subscriber added does not initially display its Distributions.

3 Select an individual Subscriber to expand its Distribution list.

or

Click *Expand All* to display the Distributions for each displayed Subscriber.

Icons indicate the status of the Subscribers' Distributions:

| Icon | Meaning |
|---|---|
|  | The Distribution was not successfully received, or the extraction failed. Check the Subscriber's event log for an error message describing the problem. See Section 2.4.2, "Managing the Distributor Agent," on page 77 . |
|  | The Distribution's status is not yet known, because the Distributor could not be contacted, or because the Distributor has not yet received the status from the Subscriber. Check the status of the Subscriber or the Distributor. See Section 2.4.2, "Managing the Distributor Agent," on page 77 . |
|  | The Distribution was successfully received, but for a non-critical reason it has not yet been extracted, such as the Extract schedule has not yet started. Check the Subscriber's extract schedule to see whether extraction has been attempted. |
|  | The Distribution has been successfully received and extracted. |

The Subscribers and Distributions are sorted by status, then alphabetically within a status. To display the more critical Distribution statuses first, the status order is:

1. Critical
2. Unknown
3. Received
4. Extracted

When one of these icons appear next to:

- ♦ **“Subscriber” (root item in tree structure):** The status applies to one or more of the subordinate Distributions. Therefore, “Subscribers” shows the most critical icon of any status in the list.
- ♦ **Subscriber icon:** The status applies to one or more of the Subscriber's Distributions. Therefore, the status icon shows the most critical status icon for any of the Subscriber's Distributions.
- ♦ **Distribution icon:** The status only applies to this Distribution.

4 Mouse-over a Distributions to display the following information:

DNS Name
NDS Name
TED Version
Receive Status

Time Received

Should Extract (either True or False displays to indicate whether the Distribution is subscribed to by the Subscriber)

Time Extracted

Extraction Status

Distributor

Parent Subscriber

If you mouse-over a status icon, it displays a short sentence of the most critical status for the related object (Subscriber or Distribution).

- 5 Fill in the time space in the *Refresh screen every __ seconds* field, then click *Start* to refresh the display at that frequency.

Only seconds can be entered.

Click *Stop* to discontinue refreshing.

This is useful for determining whether a correction to a Distribution worked, or to troubleshoot the distribution process as it rolls out to different Subscribers.

- 6 To display status information, select a Subscriber, then click *Remote Web Console*.

This option does not apply to Distributions.

For information about the types of status information you can obtain, see [Section 2.4, “Managing the Agents through Remote Web Console,” on page 72](#).

- 7 To edit configuration information, select a Subscriber or Distribution, then click *eDirectory Configuration*.

You can edit the Subscriber or Distribution object properties just as if you had clicked *Edit TED Object* under the *ZENworks Server Management* role. The same property pages and options are available in iManager that are available in ConsoleOne.

2.4 Managing the Agents through Remote Web Console

On NetWare[®] servers, you can monitor the Distributor Agent and the Policy/Package Agent at the server console where they are running. In addition, you can monitor the agents running on any supported platform (NetWare, Windows, Linux, or Solaris) from Internet Explorer using the ZENworks Server Management role in iManager.

- ♦ [Section 2.4.1, “Setting Up Passwords for Remote Web Console,” on page 73](#)
- ♦ [Section 2.4.2, “Managing the Distributor Agent,” on page 77](#)
- ♦ [Section 2.4.3, “Managing the Policy/Package Agent,” on page 80](#)
- ♦ [Section 2.4.4, “Opening Multiple Remote Web Console Windows,” on page 81](#)

However, the following Policy and Distribution Services management tasks cannot be performed in iManager and must be performed using ConsoleOne:

- ♦ Creating, editing, and deleting policy packages. See [Chapter 4, “Server Policies,” on page 193](#).
- ♦ Creating, editing, and deleting software packages. See [Chapter 5, “Server Software Packages,” on page 237](#).

2.4.1 Setting Up Passwords for Remote Web Console

To secure the features provided by Remote Web Console, you can add a password in one of the following ways:

- ♦ “Adding a Password Using iManager” on page 73
- ♦ “Adding a Password by Editing the Zws.properties File” on page 74
- ♦ “Adding a Password Using a Distributed Server Package” on page 74
- ♦ “Removing Password Protection Using iManager” on page 75

Adding a Password Using iManager

- 1 Open **iManager** and click *ZENworks Server Management > Remote Web Console*.
- 2 In the *Display* field, select *Policy Package Agent*.
- 3 Click the *Actions* tab, then click *Set Password*.

The following is displayed:

The screenshot shows the ZENworks Server Management Web Console interface. At the top, it says "ZENworks Server Management Web Console" with a help icon. Below that, the server name is "distributor-1nw.provo.novell.com" and the display is set to "Policy Package Agent". There are links for "Detach" and "View Services". The "Actions" tab is selected, and the "Set Password" option is highlighted. Below the tabs, there is a message: "Setting a password for this server adds additional access control for Web Console. Setting the password to a blank value will remove password protection." There are three input fields: "Old Password:", "New Password:", and "Confirm New Password:". An "OK" button is at the bottom.

- 4 If a previous password exists, then enter it in the *Old Password* field; otherwise, leave the field blank.

IMPORTANT: Passwords are case sensitive.

- 5 Enter the new password twice, once in the *New Password* field and again in the *Confirm New Password* field, then click *OK*.

The following is displayed:

The screenshot shows the ZENworks Server Management Web Console interface. At the top, the title is "ZENworks Server Management Web Console" with a help icon. Below the title, the server name is "distributor-1nw.provo.novell.com" and the display name is "Policy Package Agent". There are links for "Detach" and "View Services". A navigation bar contains tabs for "Configuration", "Policies", "Software Packages", "Schedule", and "Actions". Below the navigation bar, there are links for "Down Server", "Refresh", and "Set Password". The main content area contains a message: "Setting a password for this server adds additional access control for Web Console. Setting the password to a blank value will remove password protection." Below this is a confirmation message: "The password has be successfully set." There are three input fields: "Old Password:", "New Password:", and "Confirm New Password:". At the bottom, there is an "OK" button.

You do not need to click *OK* again. Clicking *OK* does not exit the page; it only causes the entries in these two fields to be validated.

The password is requested the next time Remote Web Console is accessed, even without reopening iManager.

Adding a Password by Editing the `Zws.properties` File

- 1 Open `zws.properties` in a text editor that is appropriate for the following platforms:

Linux: `/etc/opt/novell/zenworks/zws.properties`

NetWare: `volume:\zenworks\zws\zws.properties`

Windows: `drive:\zenworks\zws\zws.properties`

- 2 Locate or add the following line:

```
xmlrpcPassword=
```

This line is case sensitive.

- 3 Either replace the old password or append your new password to this line.

The password is case sensitive.

- 4 Save and exit the `zws.properties` file.

The password is requested the next time Remote Web Console is accessed, even without reopening iManager.

Adding a Password Using a Distributed Server Package

- 1 In ConsoleOne, create a Distributed Server Package.

For more information, see [Section 4.2.2, “Creating a Policy Package Object,” on page 205](#).

Depending on the purpose of the policy package, provide a descriptive package name, such as “Distributed Server Package - New RWC Password” or “Distributed Server Package - Replace RWC Password.”

2 Right-click the newly created Distributed Server Package object, then click *Properties*.

3 On the *Policies* tab, select the applicable platform, then click *Add*.

4 Select *Text File Changes*, type a name for the policy in the *Policy Name* field, then click *OK*.

The new policy should be displayed and selected. If not, select the check box in the *Enabled* column for the new policy.

5 Click *Properties*, then click *Add*.

6 In the New Text File Change dialog box, fill in the fields:

Filename: The name of the file to be edited by the policy, including its full path.

Change Description: Provide a short description of the change.

7 Depending on whether you are replacing an existing password or inserting the password for the first time, do one of the following:

- ♦ If you are creating a new password, fill in the fields:

Change Mode: Select *Append to File* from the drop-down list.

New String: Enter `xmlrpcPassword=your_new_password`, which provides this new line and password. Both the key and password are case sensitive.

- ♦ If you are replacing an existing password, fill in the fields:

Change Mode: Select *Search File* from the drop-down list.

Search Type: Select *Start of Line* from the drop-down list.

Search String: Type the search string (the beginning of the line).

Case Sensitive: Select this check box to enable it.

Find All Occurrences: Select this check box to enable it.

Result Action: Select *Replace Line* from the drop-down list.

New String: Enter `xmlrpcPassword=your_new_password`, which fully replaces the existing line and password. Both the key and password are case sensitive.

8 For multiple platforms, repeat **Step 3** through **Step 7**.

9 Click *OK* to save the changes.

10 Create a Distribution for this package and assign it to the Distributor.

For more information, see [Section 3.4.4, “Creating a Distribution,” on page 123](#).

The password is requested the next time Remote Web Console is accessed after this Distribution has been applied, even without reopening iManager.

Removing Password Protection Using iManager

1 **Open iManager** and click *ZENworks Server Management > Remote Web Console*.

2 In the *Display* field, select *Policy Package Agent*.

3 Click the *Actions* tab, then click *Set Password*.

The following is displayed:

The screenshot shows the ZENworks Server Management Web Console interface. At the top, the title is "ZENworks Server Management Web Console" with a help icon. Below the title, the server name is "distributor-1nw.provo.novell.com" and the display name is "Policy Package Agent". There are links for "Detach" and "View Services". A navigation bar contains tabs for "Configuration", "Policies", "Software Packages", "Schedule", and "Actions". Below the navigation bar, there are links for "Down Server", "Refresh", and "Set Password". The main content area contains the following text: "Setting a password for this server adds additional access control for Web Console. Setting the password to a blank value will remove password protection." Below this text are three input fields: "Old Password:", "New Password:", and "Confirm New Password:". At the bottom of the dialog is an "OK" button.

- 4 In the *Old Password* field, enter the current password.

IMPORTANT: Passwords are case sensitive.

- 5 Make sure that both the *New Password* field and the *Confirm New Password* field are empty, then click *OK*.

The following is displayed:

The screenshot shows the ZENworks Server Management Web Console interface, similar to the previous one. The main content area now contains the following text: "Setting a password for this server adds additional access control for Web Console. Setting the password to a blank value will remove password protection." Below this text is a message: "The password has be successfully set." Below the message are three input fields: "Old Password:", "New Password:", and "Confirm New Password:". At the bottom of the dialog is an "OK" button.

You do not need to click *OK* again. Clicking *OK* does not exit the page; it only causes the entries in these two fields to be validated.

The password is no longer requested the next time Remote Web Console is accessed, even without reopening iManager.

2.4.2 Managing the Distributor Agent

To access the Remote Web Console for a Distributor or Subscriber:

- 1 Open **iManager**, select *Roles and Tasks* in the top frame, expand *ZENworks Server Management* in the left frame, then click *Remote Web Console* to display the following:

ZENworks Server Management Web Console

IP address or DNS name:

Distributor, Subscriber, or Server object name:
  

- 2 Specify the IP address or DNS hostname of a server where the Distributor Agent or Policy/Package Agent is running, then click *OK*.

or

Browse to and select a Distributor or Subscriber object or the Server object representing the server where the Distributor Agent is running, then click *OK*.

If you have **passwords in effect**, the following is displayed:

ZENworks Server Management Web Console 

This server is password protected. Please enter the password to access Remote Web Console.

Password:

- 3 Enter a valid Remote Web Console password.
The password is case sensitive.
- 4 Click *Help* on each Remote Web Console page for information on using the features available on that page.

In the *Display* field, *Tiered Electronic Distribution* is the default. The other option is *Policy Package Agent* (see [Section 2.4.3, “Managing the Policy/Package Agent,”](#) on page 80).

Tabs at the top of the *Remote Web Console* frame provide various types of information related to the Policy and Distribution Services agents. Additional options are available on each tab.



- 5 Continue with the task that you want to perform:
 - ♦ [“Managing Tiered Electronic Distribution Objects”](#) on page 78
 - ♦ [“Monitoring Policy and Distribution Services Agent Status”](#) on page 78
 - ♦ [“Monitoring Distribution Status”](#) on page 79

- ◆ “Forcing Policy and Distribution Services Agent Actions” on page 79
- ◆ “Managing Security Certificates” on page 79

Table 2-1 through Table 2-5 summarize these tasks, give details for the Remote Web Console tab and option to use for each task, and indicate whether the task can also be performed using ConsoleOne.

Managing Tiered Electronic Distribution Objects

Table 2-1 *Policy and Distribution Services Agent Tasks*

| Policy and Distribution Services Agent Management Task | Remote Web Console Tab and Option | ConsoleOne |
|--|---|------------|
| List all object properties of Distributor and Subscriber objects in a single list | <i>Configuration > Configuration</i> | No |
| List the object properties of any subordinate Subscriber in the routing hierarchy | <i>Configuration > Subordinate Configuration</i> | No |
| List all object properties of Distribution objects (except type-specific information) in a single list | <i>Distributions > Distribution Information</i> | No |
| List all object properties of Channel objects in a single list | <i>Channels > Channel Information</i> | No |
| Display information about the Server Management database | <i>Configuration > Database</i> | Yes |

If the Distributor has not been refreshed since changes were made to object properties in eDirectory, the object properties displayed in the Remote Web Console are different from the object properties displayed in ConsoleOne. The Remote Web Console displays object information from the point of view of the Distributor Agent.

Monitoring Policy and Distribution Services Agent Status

Table 2-2 *Monitoring Agent Status Tasks*

| Policy and Distribution Services Agent Management Task | Remote Web Console Tab and Option | ConsoleOne |
|--|--|------------|
| View and continuously refresh the current Distributor event log, complete with message severity levels | <i>Events > Distributor Event Log</i> | No |
| View and continuously refresh the current Subscriber event log, complete with message severity levels | <i>Events > Subscriber Event Log</i> | No |
| Display the current status of the various distribution threads started by the Policy and Distribution Services agents to perform their various functions | <i>Configuration > Threads</i> | No |

Monitoring Distribution Status

Table 2-3 *Monitoring Distribution Status Tasks*

| Policy and Distribution Services Agent Management Task | Remote Web Console Tab and Option | ConsoleOne |
|---|--|------------|
| List all Distributions currently being processed by the Distributor or Subscriber, along with detailed status information | <i>Distributions > Active Distributions</i> | No |
| Display status information for a selected Distribution that has been received by a Subscriber | <i>Distributions > Received Distributions</i> | No |
| Display the route that a Distribution must take through the routing hierarchy from a Distributor or parent Subscriber to any subordinate Subscriber | <i>Configuration > Route to Subscriber</i> | No |

Forcing Policy and Distribution Services Agent Actions

Table 2-4 *Forcing Agent Actions Tasks*

| Policy and Distribution Services Agent Management Task | Remote Web Console Tab and Option | ConsoleOne |
|---|---|--------------------|
| Immediately refresh a Distributor so that it reads eDirectory to check for modified Distributions | <i>Configuration > Refresh Distributor</i> | Yes |
| Immediately build a Distribution | <i>Distributions > Build Distribution</i> | Schedule dependent |
| Immediately send all Distributions listed in a selected Channel | <i>Channels > Distribute Channel</i> | Not with one click |

Managing Security Certificates

Table 2-5 *Managing Security Certificates Tasks*

| Policy and Distribution Services Agent Management Task | Remote Web Console Tab and Option | ConsoleOne |
|--|--|------------|
| List the security certificates that are available on a Subscriber | <i>Security > Show Certificates</i> | No |
| Delete security certificates from a Subscriber | <i>Security > Show Certificates > Remove Certificate</i> | No |
| Have the Distributor sign Subscribers' Certificate Signing Request (.CSR) files so that the Subscribers can receive encrypted Distributions from the Distributor | <i>Security > Sign CSR</i> | Yes |

2.4.3 Managing the Policy/Package Agent

The Policy/Package Agent is responsible for installing the software and enforcing the policies that it receives and extracts. The Remote Web Console enables you to manage the Policy/Package Agent, which is not possible using ConsoleOne.

To access the Remote Web Console for a Policy/Package Agent:

- 1 **Open iManager**, select *Roles and Tasks* in the top frame, expand *ZENworks Server Management* in the left frame, then click *Remote Web Console* to display the following:

ZENworks Server Management Web Console

IP address or DNS name:

Distributor, Subscriber, or Server object name:
  

- 2 Specify the IP address or DNS hostname of a server where the Policy/Package Agent is running, then click *OK*.

or

Browse to and select a Subscriber object or the Server object representing the server where the Policy/Package Agent is running, then click *OK*.

- 3 If you have **passwords in effect**, the following is displayed:

ZENworks Server Management Web Console 

This server is password protected. Please enter the password to access Remote Web Console.

Password:

Enter a valid Remote Web Console password.

The password is case sensitive.

- 4 In the *Display* field, select *Policy Package Agent* (*Tiered Electronic Distribution* is the default). Click *Help* on each Remote Web Console page for information on using the features available on that page.

Tabs at the top of the *Remote Web Console* frame provide various types of information related to the Policy/Package Agent.



- 5 Continue with the task that you want to perform.

The following table summarizes these tasks and gives the *Remote Web Console* tab for each task.

| Policy/Package Agent Management Task | Remote Web Console | ConsoleOne |
|--|---------------------------|-------------------|
| List the plug-ins that are currently loaded for enforcing server policies | Configuration | No |
| List all the variables that the Policy/Package Agent has values for | Configuration | No |
| List all the policies that the Policy/Package Agent enforces on a Subscriber server | Policies | No |
| Immediately enforce one or more policies on a Subscriber server | Policies | No |
| Remove individual policies from a Subscriber server | Policies | No |
| Immediately refresh one or more policies so that the Distributor Agent reads eDirectory to check for modifications | Policies | No |
| List all the software packages that the Policy/Package Agent installs on the Subscriber server | Software Packages | No |
| Determine the current status of all software packages installed on the Subscriber server | Software Packages | No |
| Create and run a program or script on the Subscriber server once or repeatedly | Schedule | No |
| Down the Subscriber server | Actions | No |
| Restart the Policy/Package Agent | Actions | No |

2.4.4 Opening Multiple Remote Web Console Windows

On any Remote Web Console page, click *Detach* in the upper right corner to display the current page in a new browser window. This enables you to access multiple Remote Web Console features at the same time. For example, you could detach one window for the Tiered Electronic Distribution agents and another window for the Policy/Package Agent. Or you could detach a window for the Remote Web Console and still be able to perform other ZENworks Server Management tasks in the main Novell iManager window.

2.5 Comparing the ZENworks Server Management Role in iManager with ConsoleOne Capabilities

Table 2-6 Differences between iManager and ConsoleOne

| Task | iManager | ConsoleOne |
|--|----------|--------------------|
| Creating, editing, and deleting the following Tiered Electronic Distribution objects: Distributor (cannot create) Subscriber (cannot create) Distribution Channel Subscriber Group External Subscriber | Yes | Yes |
| Creating, editing, and deleting the following Policy and Distribution Services components: Policy Package Server Software Package Desktop Application | No | Yes |
| Setting up the following Distribution types: Desktop Application File FTP HTTP MSI Policy Package RPM Software Package | Yes | Yes |
| Immediately refreshing a Distributor | Yes | Yes |
| Immediately building a Distribution | Yes | Not with one click |
| Immediately sending to Subscribers all Distributions listed in a Channel | Yes | Not with one click |
| Monitoring Policy and Distribution Services agent event logs and status | Yes | No |
| Listing and managing the policies on a Subscriber server | Yes | No |
| Listing and checking the status of software packages installed on a Subscriber server | Yes | No |
| Running programs and scripts on a Subscriber server | Yes | No |
| Downing a Subscriber server | Yes | No |

| Task | iManager | ConsoleOne |
|--|-----------------|-------------------|
| Managing security certificates: | | |
| Listing available certificates | Yes | No |
| Resolving certificates | No | Yes |
| Signing CSRs | Yes | Yes |
| Managing the Policy/Package Agent | Yes | No |

Tiered Electronic Distribution

3

Novell® ZENworks® Server Management provides Tiered Electronic Distribution for managing distributions of files, policies, and software across your network.

Tiered Electronic Distribution is integrated with other Novell network management applications that snap in to the ConsoleOne® framework to take advantage of Novell eDirectory™ management and file access control. Tiered Electronic Distribution can also be managed using the ZENworks Server Management role in Novell iManager.

For information on Tiered Electronic Distribution, see the following sections:

- ◆ [Section 3.1, “Common Distribution Tasks,” on page 85](#)
- ◆ [Section 3.2, “Understanding Tiered Electronic Distribution,” on page 87](#)
- ◆ [Section 3.3, “Distributors,” on page 95](#)
- ◆ [Section 3.4, “Distributions,” on page 110](#)
- ◆ [Section 3.5, “Channels,” on page 144](#)
- ◆ [Section 3.6, “Subscribers,” on page 147](#)
- ◆ [Section 3.7, “Subscriber Groups,” on page 155](#)
- ◆ [Section 3.8, “External Subscribers,” on page 156](#)
- ◆ [Section 3.9, “Configuring Multiple Tiered Electronic Distribution Objects,” on page 165](#)
- ◆ [Section 3.10, “Sending Distributions,” on page 172](#)
- ◆ [Section 3.11, “Miscellaneous Tiered Electronic Distribution Issues,” on page 175](#)
- ◆ [Section 3.12, “Working Directories,” on page 186](#)
- ◆ [Section 3.13, “Editing the Tednode.properties File,” on page 190](#)

3.1 Common Distribution Tasks

[Table 3-1](#) through [Table 3-6](#) provide documentation links to common Tiered Electronic Distribution tasks. All links are to sections in this Policy and Distribution Services portion of the *Administration* guide.

Tiered Electronic Distribution Objects

Table 3-1 *Common Tiered Electronic Distribution Tasks*

| Task | Instructions |
|--|--|
| Create a Distributor or Subscriber | “Installation on NetWare and Windows Servers” in the <i>Novell ZENworks 7 Server Management Installation Guide</i> |
| Configure multiple Tiered Electronic Distribution objects | Section 3.9, “Configuring Multiple Tiered Electronic Distribution Objects,” on page 165 |
| Change the DNS name or IP address of a Tiered Electronic Distribution server | “Changing DNS Names or IP Addresses for Tiered Electronic Distribution Servers” on page 185 |

Distributor

Table 3-2 *Common Distributor Tasks*

| Task | Instructions |
|--|---|
| Configure a Distributor object | "Configuring Distributors" on page 106 |
| Create a routing hierarchy for a Distributor | "Understanding Distribution Routing" on page 97 and "Configuring Distributors" on page 106 |
| Delete a Distributor object | "Deleting a Distributor Object and How Its Distributions Are Affected" on page 110 |
| Refresh a Distributor | "Manually Refreshing the Distributor" on page 109 |
| Create a security certificate on a Distributor and copy it to its associated Subscribers | "Creating Security Certificates for Non-Encrypted Distributions" on page 307 |

Distribution

Table 3-3 *Common Distribution Tasks*

| Task | Instructions |
|--|--|
| Create a Distribution | Section 3.4, "Distributions," on page 110 |
| Delete a Distribution | "Deleting a Distribution" on page 136 |
| Managing orphaned Distributions (when their Distributor object has been deleted) | "Deleting a Distributor Object and How Its Distributions Are Affected" on page 110 |
| Schedule and send a Distribution | Section 3.10, "Sending Distributions," on page 172 |
| Force a Distribution to be sent | "Forcing a Single Distribution To Be Sent" on page 173 |
| Use a parent Subscriber to send a Distribution | "Sending Distributions Through Parent Subscribers" on page 174 |
| Send a Distribution to another tree | "Sending Distributions between Trees" on page 174 |
| Import or export a Distribution manually | "Manually Importing and Exporting Distributions" on page 141 |
| Create and send a File Distribution using a wizard | "Using the Distribution Wizard" on page 143 |

Channel

Table 3-4 *Common Channel Tasks*

| Task | Instructions |
|------------------|---|
| Create a Channel | "Creating and Configuring Channels" on page 146 |

| Task | Instructions |
|-------------------------|--|
| Force a Channel to fire | “Forcing a Channel To Be Sent” on page 147 |

Subscriber

Table 3-5 *Common Subscriber Tasks*

| Task | Instructions |
|--|--|
| Configure a Subscriber object | “Configuring Subscribers” on page 150 |
| Create an External Subscriber object | “Creating and Configuring External Subscribers” on page 164 |
| Configure the <code>tednode.properties</code> file for a Subscriber server that does not have its own configuration capability | Section 3.13, “Editing the Tednode.properties File,” on page 190 |

Network Traffic Management

Table 3-6 *Common Network Traffic Management Tasks*

| Task | Instructions |
|---|--|
| Control bandwidth usage for Distribution traffic by setting the I/O rates | “Controlling I/O Rates and Concurrent Distributions” on page 182 |
| Minimize network messaging traffic | “Minimizing Messaging Traffic” on page 183 |

3.2 Understanding Tiered Electronic Distribution

Review the following sections for an understanding of Tiered Electronic Distribution:

- ◆ [Section 3.2.1, “Distribution Management through Tiered Electronic Distribution,” on page 88](#)
- ◆ [Section 3.2.2, “The Basic Distribution Process,” on page 88](#)
- ◆ [Section 3.2.3, “Tiered Electronic Distribution’s eDirectory Objects,” on page 89](#)
- ◆ [Section 3.2.4, “Relationships of the Tiered Electronic Distribution Objects,” on page 89](#)
- ◆ [Section 3.2.5, “Physical Network Connections,” on page 90](#)
- ◆ [Section 3.2.6, “Distribution Flow Details,” on page 90](#)
- ◆ [Section 3.2.7, “Tiered Electronic Distribution Processes,” on page 91](#)
- ◆ [Section 3.2.8, “The Tiered Distribution Model,” on page 93](#)
- ◆ [Section 3.2.9, “Tiered Electronic Distribution’s Key Components,” on page 94](#)

3.2.1 Distribution Management through Tiered Electronic Distribution

Tiered Electronic Distribution provides you with a way to manage your servers through the distribution of electronic data between servers. For example, application programs, collections of data files, software patches, and server policies.

When you install Policy and Distribution Services, the installation process creates Tiered Electronic Distribution and server policy objects in the eDirectory tree, copies software to the various servers, and sets up basic configurations for the Tiered Electronic Distribution and Server Policies components according to your installation selections.

The Tiered Electronic Distribution software can be hosted on NetWare[®], Windows 2000, Windows 2003 Server, Linux, and Solaris servers.

Tiered Electronic Distribution uses a tiered distribution model that enables one server to indirectly service hundreds or even thousands of other servers. Tiered Electronic Distribution makes it easy to distribute files and policy packages by building them into compressed data files and hosting them in distribution channels for dissemination to the appropriate servers.

Tiered Electronic Distribution lets you schedule the distribution processes to take advantage of off-peak hours. It also sends notification of distribution status by sending e-mail messages, logging events, displaying real-time messages, database reporting, and sending SNMP traps.

Server Management can efficiently process (send/receive/extract) Distributions that are large in size and contain a substantial number of files, such as an entire 4GB volume with greater than 50,000 file entries.

3.2.2 The Basic Distribution Process

The Tiered Electronic Distribution distribution process is based on the creation of Distributions (compressed file collections) that you use to move files and policies to your network servers. For more information, see [Section 3.10.1, “Understanding the Distribution Processes,” on page 172](#).

Following is a simplified distribution process. It is governed by [schedules](#) that you set for each of the Tiered Electronic Distribution objects involved with the Distribution file.

1. A Distributor creates a [security certificate](#) to provide distribution security.
2. A Distribution is built on the Distributor server’s file system according to the configuration you create in the [Distribution object](#).
3. You associate the Distribution with a [Channel](#).
4. You [subscribe](#) your target [Subscriber servers](#) to the Channel. This causes them to receive all of the Distributions contained in that Channel.
5. The certificate (from 1 above) is copied to Subscriber servers for Distribution security verification.
6. The Channel’s listed Distributions are sent from the Distributor to the Subscriber servers whose security certificates are valid.
7. The Subscriber extracts the files or policies from the compressed Distribution file and applies them according to the Distribution object’s configuration.

The schedules that you need to coordinate for sending Distributions are the Distributor’s Refresh schedule, the Distribution’s Build schedule, and the Channel’s Send schedule. However, we recommend that you leave the Distributor’s Refresh schedule set to the default of Never. For more information, see [“Determining the Distributor’s Refresh Schedule” on page 49](#).

The schedules that you need to coordinate for receiving and extracting Distributions are the Channel’s Send schedule and the Subscriber’s Extract schedule.

3.2.3 Tiered Electronic Distribution’s eDirectory Objects

Tiered Electronic Distribution uses eDirectory objects and the related software for performing its distribution functions. The Distinguished Name (DN) of all Tiered Electronic Distribution objects includes the server name and component function of the host server.

The eDirectory schema extensions included in Tiered Electronic Distribution define the classes of eDirectory objects that are created in your eDirectory tree, including information that is required or optional at the time the object is created. Every object associated with Tiered Electronic Distribution in an eDirectory tree has a class defined for it in the tree’s schema.

You will extend the schema of your tree for the eDirectory objects listed in [Table 3-7](#) when you install ZENworks 7 Server Management:

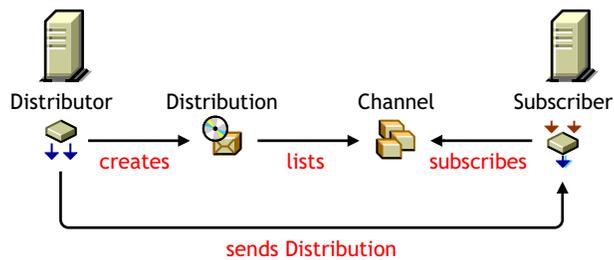
Table 3-7 Tiered Electronic Distribution eDirectory Objects

| Object | Basic Function | More Information |
|---------------------|---|--|
| Distributor | Build, send Distributions | Section 3.3, “Distributors,” on page 95 |
| Distribution | Contain files, policies | Section 3.4, “Distributions,” on page 110 |
| Channel | List Distributions | Section 3.5, “Channels,” on page 144 |
| Subscriber | Receive, extract Distributions | Section 3.6, “Subscribers,” on page 147 |
| Subscriber Group | Channel subscriptions by multiple Subscribers | Section 3.7, “Subscriber Groups,” on page 155 |
| External Subscriber | Enable distributing between trees | Section 3.8, “External Subscribers,” on page 156 |

3.2.4 Relationships of the Tiered Electronic Distribution Objects

[Figure 3-1](#) illustrates the relationships of the main Tiered Electronic Distribution objects:

Figure 3-1 *The Distributor, Distribution, Channel, Subscriber, and External Subscriber Objects*



Note the following from this illustration:

- ◆ A Distributor creates a Distribution
- ◆ The Distribution is listed in a Channel
- ◆ A Subscriber subscribes to the Channel
- ◆ The Subscriber receives the Distribution from the Distributor (possibly via a parent Subscriber)

3.2.5 Physical Network Connections

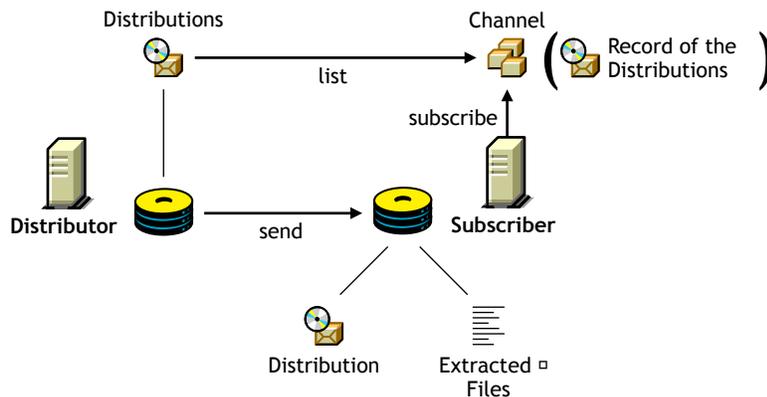
Distributor and Subscriber servers can be physically connected to the network in any configuration, including having some servers across WAN links. The following describes the possible physical interactions between Distributor and Subscriber servers:

- ◆ A Subscriber server can be in the same geographic location as its Distributor server
- ◆ A Subscriber server can be in a different geographic location from its Distributor server, such as across a WAN link
- ◆ A Distributor server can service multiple Subscriber servers
- ◆ A Subscriber server can be serviced by multiple Distributor servers
- ◆ A Subscriber server can receive its Distribution files directly from a Distributor server
- ◆ A Subscriber server can receive its Distribution files indirectly via another Subscriber server acting as a parent Subscriber

3.2.6 Distribution Flow Details

Figure 3-2 illustrates the physical flow of Tiered Electronic Distribution Distributions:

Figure 3-2 Tiered Electronic Distribution Flow



Note the following from the illustration:

- ◆ A Distribution file is stored on the Distributor server's hard drive
- ◆ The Channel lists a Distribution (it does not hold a copy of the Distribution)
- ◆ The Subscriber subscribes to a Channel to obtain all of the Distributions listed there
- ◆ The Subscriber extracts the Distribution contents from the file's compressed format and writes the content to the volume and directory specified in the Distribution's configuration

IMPORTANT: When there are multiple versions of a File or Desktop Application Distribution, the Subscriber maintains copies of each of the versions, as is specified in the Distribution object's properties. The default is to maintain 10 versions per Distribution type.

3.2.7 Tiered Electronic Distribution Processes

The following processes are used to perform Tiered Electronic Distribution functions:

- ◆ ["Distributor Agent" on page 91](#)
- ◆ ["Policy/Package Agent" on page 92](#)
- ◆ ["Tiered Electronic Distribution Software Running on the Subscriber Server" on page 92](#)
- ◆ ["Distribution Processes Summary" on page 93](#)

Distributor Agent

The Distributor Agent is installed on each server where you select the Distributor option during installation.

This agent has the following functions:

- ◆ Reads eDirectory for all Tiered Electronic Distribution configuration information (Distribution, Channel, and Subscriber) according to the Refresh schedule
- ◆ Builds Distributions based on the information contained in the Distribution objects that are associated with the Distributor
- ◆ Builds Distributions according to the Build schedule
- ◆ Sends Distributions according to the Send schedule

- ◆ Handles all notifications and events for the Subscriber
- ◆ Sends DS configuration information found in Subscriber objects to each Subscriber as part of each Distribution
- ◆ Logs Tiered Electronic Distribution information to the `ted.log` file for reporting purposes

Policy/Package Agent

The Policy/Package Agent is installed on each server where you selected the Policy and Distribution Server option during installation.

This agent has the following Tiered Electronic Distribution functions:

- ◆ Reads and enforces policy information that has been extracted from Policy Package Distributions

For more information on policies, see [Chapter 4, “Server Policies,” on page 193](#).

- ◆ Installs Server Software Packages that have been extracted from Software Package Distributions

For more information on software packages, see [Chapter 5, “Server Software Packages,” on page 237](#).

- ◆ Logs policy and software package information to the `zfs-startup.log` file for reporting purposes

Tiered Electronic Distribution Software Running on the Subscriber Server

Tiered Electronic Distribution software is installed on each server where you selected the Policy and Distribution Server option during installation.

This software has the following functions:

- ◆ Subscribes a Subscriber server to Channels for receiving Distributions
- ◆ Receives and extracts the following Distribution types to the server’s file system according to the Extract schedule:

Desktop Application ¹

File

FTP

HTTP

MSI

Policy Package

RPM

Software Package

¹ The Desktop Application Distribution is only available when ZENworks Desktop Management is installed.

- ◆ Installs the following extracted Distributions:

Desktop Application

MSI

RPM

- ◆ In the parent Subscriber role, receives a Distribution and forwards it on to other Subscriber servers

Distribution Processes Summary

Table 3-8 *The Distribution Processes*

| Function | Process | Explanation |
|------------------------------------|--|--|
| Building and Sending Distributions | Distributor Agent | Discovers, builds, and sends all Distributions using the Distributor server's CPU and file system. |
| Extracting Distributions | Tiered Electronic Distribution software running on the Subscriber server | Extracts the Distribution's data onto the Subscriber server using the Subscriber server's CPU and file system. Also notifies the Policy/Package Agent when there are Server Policies to be enforced, or Server Software Packages to be installed. |
| Installing Distributed Software | Policy/Package Agent | Installs Server Software Packages onto the Subscriber server using the Subscriber server's CPU and file system. |
| Enforcing Installed Policies | Policy/Package Agent | Reads and enforces the extracted policies on the Subscriber server using its CPU and file system. |

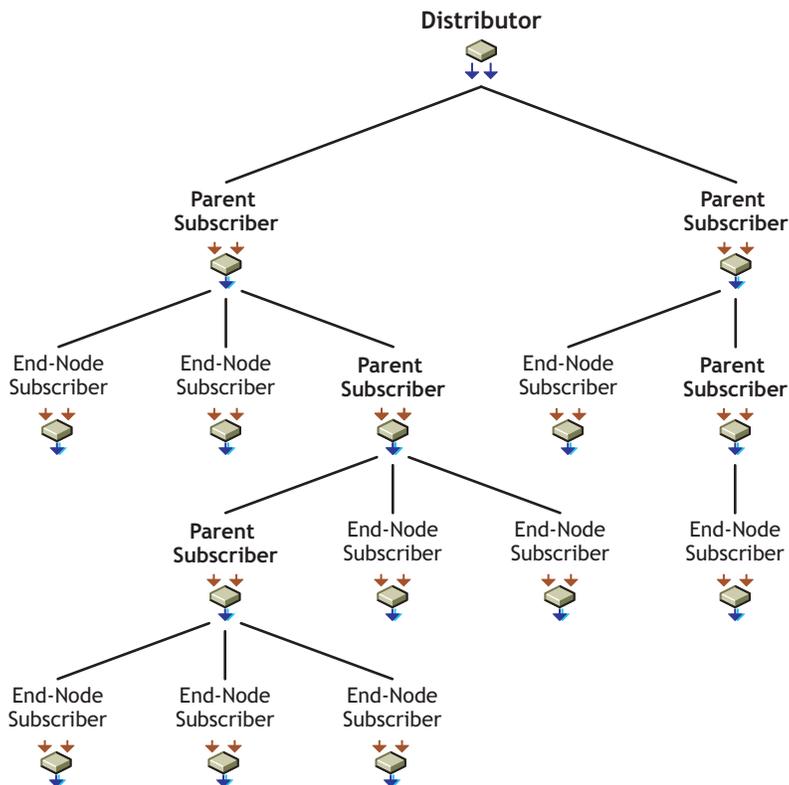
3.2.8 The Tiered Distribution Model

The power of the tiered distribution model is that you can spread the workload for sending Distributions. This is particularly important to the Distributor servers. By sharing distribution duties with parent Subscribers, a Distributor server can have more resources available for reading eDirectory, building each of its Distributions, and logging information to the database.

Tiered distribution levels can be very deep, providing a very large number of Subscribers that any one Distributor can service—without doing so directly.

Figure 3-3 illustrates a distribution routing hierarchy containing a Distributor, several parent Subscribers, and many end-node Subscribers:

Figure 3-3 *Distribution Route Hierarchy Showing Parent Subscribers and End-Node Subscribers*



The Distributor can service hundreds of parent Subscribers directly, or service just a few first-tier parent Subscribers and let them do the bulk of the distribution work. In the above illustration, the Distributor only sends its Distribution to two parent Subscribers, yet nine end-node Subscribers receive the Distribution.

The parent Subscribers shown in this illustration can also receive the Distribution for extraction if they were also subscribed to the Distribution's Channel. If all of the parent Subscribers in the above illustration were subscribed to receive the Distribution being sent to the end-node Subscribers, the Distributor services 14 total Subscriber servers while itself sending the Distribution only twice.

Each parent Subscriber can service hundreds of other parent Subscribers or end-node Subscribers (the intended recipients of the Distributions). The workload for passing on a Distribution by a parent Subscriber is minimal in compared to the workload for the Distributor to build the Distribution.

As you can see, the tiered distribution model allows you to minimize the distribution workload for your Distributor servers.

3.2.9 Tiered Electronic Distribution's Key Components

In summary, the key components of Tiered Electronic Distribution include:

- ◆ eDirectory schema extensions that include objects for Distributors, Distributions, Channels, Subscribers, and External Subscribers
- ◆ ConsoleOne snap-ins and iManager plug-ins that provide creation, configuration, and management of Tiered Electronic Distribution

- ◆ A Distributor Java process hosted on a NetWare, Windows 2000, Windows 2003 Server, Linux, or Solaris server for handling distribution of data packages to Subscribers
- ◆ A Subscriber Java process hosted on a NetWare, Windows 2000, Windows 2003 Server, Linux, or Solaris server that subscribes to a Channel for its Distributions
- ◆ A routing hierarchy for each Distributor that has a hierarchical list of Subscribers who can both receive Distributions for themselves and pass the Distributions on to other Subscribers
- ◆ Parent Subscribers that pass Distributions on to other Subscribers
- ◆ An External Subscriber object that allows distributing between trees or to servers that do not have eDirectory server objects
- ◆ The Distributor Agent that controls the actual processes of building the Distribution files on the Distributor
- ◆ Policy/Package Agent that extracts and enforces policy information from Policy Package Distributions, and extracts and installs the contents of software packages
- ◆ Certificates that provide distribution security

3.3 Distributors

The following sections provide concepts and instructions for the Distributor object:

- ◆ [Section 3.3.1, “Understanding Distributors,” on page 95](#)
- ◆ [Section 3.3.2, “Understanding Distribution Routing,” on page 97](#)
- ◆ [Section 3.3.3, “Creating Distributors,” on page 106](#)
- ◆ [Section 3.3.4, “Configuring Distributors,” on page 106](#)
- ◆ [Section 3.3.5, “Manually Refreshing the Distributor,” on page 109](#)
- ◆ [Section 3.3.6, “Deleting a Distributor Object and How Its Distributions Are Affected,” on page 110](#)

3.3.1 Understanding Distributors

The Distributor object (TED Distributor) is an eDirectory object that defines the properties for the Distributor.

- ◆ [“Functional Relationship with Other Tiered Electronic Distribution Objects” on page 95](#)
- ◆ [“Distributor Description” on page 96](#)
- ◆ [“Scheduling” on page 96](#)
- ◆ [“Routing Distributions” on page 97](#)
- ◆ [“Multiple Distributors in the Tree” on page 97](#)
- ◆ [“Database Logging” on page 97](#)

Functional Relationship with Other Tiered Electronic Distribution Objects

[Figure 3-4](#) illustrates that a Distributor sends its Distributions to Subscriber servers:

Figure 3-4 *Distributor Sending to Multiple Subscribers*

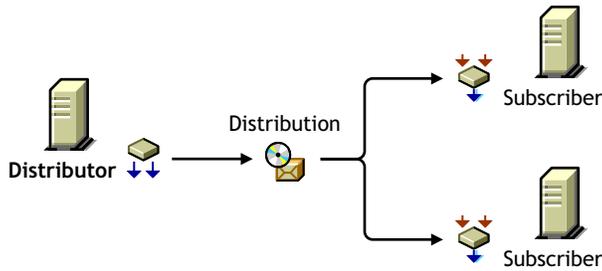
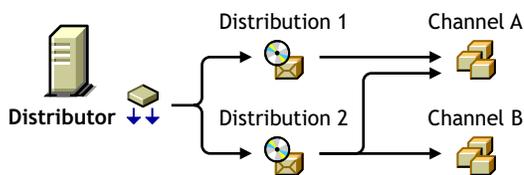


Figure 3-5 illustrates that a Distributor can list any one of its Distributions in several Channels, and several of its Distributions in one Channel:

Figure 3-5 *Distributor Listing Distributions in Multiple Channels*



Distributor Description

The Distributor server's main Tiered Electronic Distribution function is to create and send Distributions. It also logs information to a database file, if you have one assigned for the Distributor.

The Distributor Agent builds a Distribution file on the Distributor server from the information you provide when you create and configure a Distribution object. A Distributor can own multiple Distributions.

When a Distributor builds a Distribution, it can optionally create a digest that provides an MD5 checksum for the Subscriber to compare against. Digests are used by Subscribers to verify that the Distributions have not been tampered with while in transit. Creating a digest is optional per Distributor, so the digests might not always be available for a checksum comparison by any Subscriber where this option is enabled.

Digests also detect corruption in a Distribution's package. In the case of corruption, the Subscriber renames the `distfile.ted` Distribution file to `distfile.corrupt` and the Distribution is rebuilt and sent the next time the Channel's schedule fires.

A Distributor lists its Distributions in Channels. Distributors do not own Channels. However, a Distributor is the sole owner of its Distributions.

The Distributor sends its Distributions to Subscribers (usually parent Subscribers for passing on the Distributions). If an end-node Subscriber does not respond to a Distributor (or a parent Subscriber) that is trying to send a Distribution to it, the Distributor retries sending a Distribution every two minutes for 30 minutes, then stops. It does not attempt to re-send the Distribution until the Channel's Send schedule starts again.

Scheduling

A Distributor's Refresh schedule determines when it reads eDirectory for changes to its Distributions and other Tiered Electronic Distribution objects. A Distributor builds all new

Distributions it finds and rebuilds any of its Distributions that have changed. The new or rebuilt Distributions are then available to be sent when a Channel's Send schedule starts.

IMPORTANT: We recommend that the Distributor's Refresh schedule be left at the default of Never, unless you have a reason to schedule the refresh. If you set the Refresh schedule, it is possible that the Distribution building and sending processes can be interrupted and restarted. This could possibly cause an infinite loop situation where the Distributions never get built or sent.

A Distributor can build its Distributions any time its Refresh schedule starts.

If you delete a Distribution, you should also refresh the Distributor immediately so that it recognizes the deletion and not try to build a Distribution that no longer exists. For information on deleting Distributions, see [Section 3.4.8, "Deleting a Distribution," on page 136](#).

For information on scheduling, see [Chapter 8, "Scheduling," on page 317](#).

Routing Distributions

The Distributor contains a distribution route, which is a hierarchical list of Subscribers that indicate the routes the Distributor can take to send its Distributions to its Subscriber servers. For information on routing hierarchies, see ["Understanding Distribution Routes" on page 42](#).

Multiple Distributors in the Tree

You can have multiple Distributor objects in the tree; however, you can only have one Distributor installed per server. The need for multiple Distributors is dependent on several factors. For more information, see [Section 1.1.4, "Are Additional Distributors Needed?," on page 37](#).

Database Logging

Individual Distributors can log information to their own database files, or all Distributors can log information to one common database file. For information on databases, see [Chapter 10, "ZENworks Database," on page 351](#).

3.3.2 Understanding Distribution Routing

A distribution route represents the most efficient path to any given segment of your WAN. A distribution route is a list of parent Subscribers that relay Distributions on to other parent or end-node Subscribers. You can use Parent Subscribers to minimize the workload for a Distributor because they can pass on Distributions to other Subscribers.

The following sections explain how a Distributor moves its Distributions to your network's servers:

- ◆ ["Understanding Parent Subscribers" on page 98](#)
- ◆ ["Understanding Routing Hierarchies" on page 100](#)
- ◆ ["Sharing Parent Subscribers with Other Distributors" on page 102](#)
- ◆ ["Distributing Across WAN Links" on page 103](#)
- ◆ ["Out-of-Tree Distributions" on page 104](#)
- ◆ ["Routing Hierarchy Configuration Guidelines" on page 105](#)

Understanding Parent Subscribers

A parent Subscriber is a Subscriber that acts as a proxy for the Distributor to store and pass Distributions so that the Distributor does not need to send its Distributions directly to every Subscriber. Parent Subscriber servers do not need to be recipients themselves of a Distribution to temporarily store it for passing on to other Subscriber servers.

- ◆ [“Distributors Send Distributions Using Parent Subscribers” on page 98](#)
- ◆ [“Passing on Unsubscribed Distributions” on page 98](#)
- ◆ [“Sharing the Distribution Load” on page 98](#)
- ◆ [“Balancing Workloads” on page 98](#)

Distributors Send Distributions Using Parent Subscribers

A Distributor server must actually send each of its Distributions, because the Distribution files reside in its own file system.

Sending Distributions can create an enormous workload for a Distributor if it must individually send each of its Distributions to every Subscriber server on the network. Therefore, parent Subscribers are used to help send Distributions.

A detailed understanding of your network’s topology is important for properly configuring distribution routes and selecting parent Subscribers. If necessary, create a diagram of your network showing all WAN links to determine how to use parent Subscribers.

Passing on Unsubscribed Distributions

A Subscriber does not need to subscribe to a Channel containing a Distributor’s Distributions to be in the Distributor’s routing hierarchy. A parent Subscriber itself does not need to be the recipient of the Distribution it is passing on.

Further, a parent Subscriber does not need to subscribe to the same Channels as its subordinate Subscribers to be able to pass on those Channel’s Distributions.

Sharing the Distribution Load

In the illustration under [“The Routing Hierarchy” on page 100](#), each Subscriber listed could be a parent to other Subscribers on its LAN. For example, if every Subscriber listed in the illustration was a parent to 20 end-node Subscribers, the Distributor could service 210 total Subscribers while only physically sending its Distributions to three of the Subscribers (the first-tier parent Subscribers, numbers 01, 04, and 09).

To further illustrate, parent Subscriber 04 would be servicing 104 Subscribers while only directly sending to two parent Subscribers (05 and 06) and its own 20 end-node Subscribers.

Balancing Workloads

A Distributor can use parent Subscribers in a routing hierarchy to explicitly determine routes for its Distributions. This eases its workload in distributing to Subscribers.

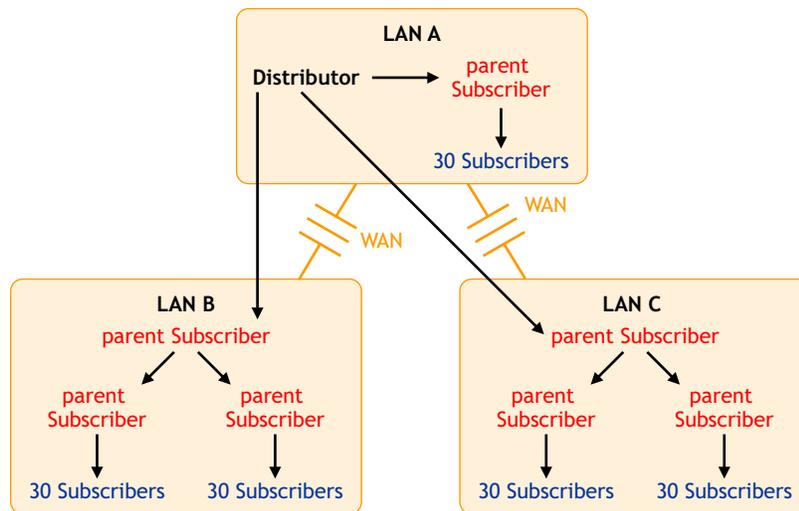
A parent Subscriber can also help a Distributor with its workload by acting as a proxy for the Distributor to pass on Distributions to other Subscribers. You can have multiple parent Subscribers on a given LAN to share the distribution workload on the LAN.

We estimate that the number of Subscribers and/or parent Subscribers that any one Distributor or parent Subscriber should service to be about 40. This figure is dependent on such factors as network speed, sizes of Distributions, and so on.

You should place parent Subscribers where they can help in load-balancing for Distributors and other parent Subscribers.

Figure 3-6 illustrates a WAN environment with parent Subscribers:

Figure 3-6 WAN Environment with Parent Subscribers



Note the following from this illustration:

- ◆ Assume that the three parent Subscribers that the Distributor's distribution lines point to are the first-tier Subscribers in the Distributor's routing hierarchy.
- ◆ Assume that the other four parent Subscribers (in LAN B and LAN C) are listed in the second tier of the distribution hierarchy.
- ◆ The Distributor does not need to send the Distributions directly to the 30 Subscribers on LAN A because the parent Subscriber in LAN A does that.
- ◆ The Distributor only sends its Distributions directly to the three parent Subscribers, but a total of 157 Subscribers can receive those Distributions.
- ◆ One parent Subscriber in LAN B (and the same for LAN C) was used solely for receiving Distributions directly from the Distributor, then passing them on to other parent Subscribers, which in turn passed them to their 60 Subscribers. For large systems, this scheme can make a parent Subscriber on the other side of a WAN link more available to a Distributor, instead of that parent Subscriber being so busy passing Distributions to its many other end-node Subscribers that it can make the Distributor wait. Consider this hierarchical design where it might be applicable in your network.

The Distributor has the workload of reading eDirectory for Distribution changes, building the Distributions, sending the Distributions, and writing to the Server Management database. By minimizing the number of Subscribers that a Distributor itself must directly send Distributions to, you can give the Distributor more resources for its various functions.

Understanding Routing Hierarchies

Tiered Electronic Distribution provides a routing hierarchy to automate sending your Distributions from the Distributor servers to your Subscriber servers.

- ◆ “The Routing Hierarchy” on page 100
- ◆ “Distributing Using the Hierarchy” on page 100
- ◆ “Subscribers Orphaned from the Routing Hierarchy” on page 101
- ◆ “Rerouting Because of Changes to the Routing Hierarchy” on page 102

The Routing Hierarchy

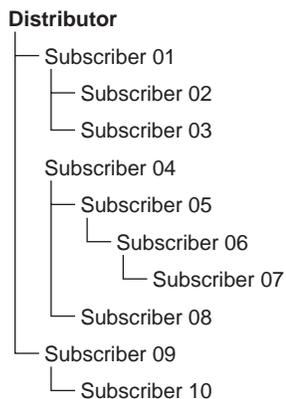
To ease a Distributor’s workload in sending Distributions, each Distributor has its own routing hierarchy, which is a hierarchical list of Subscribers that indicate the routes Distributions can take to send a Distribution to a Subscriber. The Subscribers in the routing hierarchy are the parent Subscribers. You can nest parent Subscribers many levels deep.

A parent Subscriber can receive a Distribution and extract it, as well as pass that same Distribution on to other Subscribers.

You can modify distribution routes at any time by editing the properties of the Distributor objects.

Figure 3-7 illustrates a Distributor’s routing hierarchy:

Figure 3-7 *Distributor’s Routing Hierarchy*

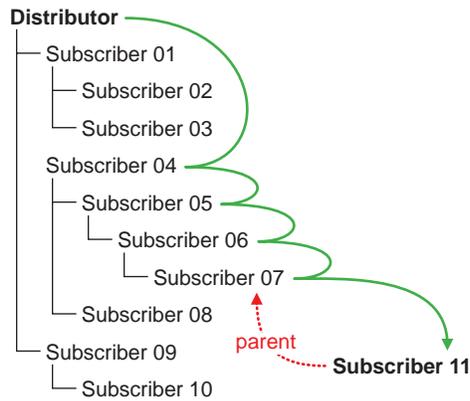


The only Subscribers you need to include in the Distributor’s routing hierarchy are those that are used to pass on Distributions to other Subscribers. Subscribers that are not used to pass on Distributions are referred to as end-node Subscribers.

Distributing Using the Hierarchy

Assume that Subscriber 07 is a parent to Subscriber 11 (which is not in the routing hierarchy). The distribution route from the Distributor to Subscriber 11 would be as shown in **Figure 3-8**:

Figure 3-8 *Distributing Within the Hierarchy*

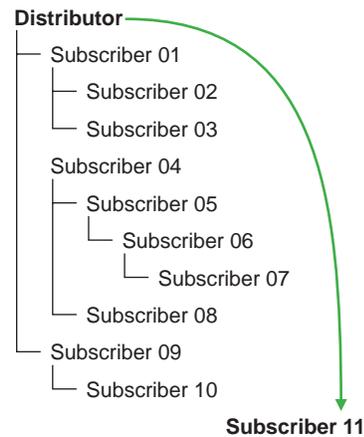


The Distributor used four parent Subscribers (04, 05, 06, and 07) to send the Distribution to Subscriber 11.

Subscribers Orphaned from the Routing Hierarchy

If Subscriber 11 did not have a parent Subscriber (such as Subscriber 07), the Distribution would come directly from the Distributor as shown in **Figure 3-9**:

Figure 3-9 *Subscribers Orphaned in the Distribution Hierarchy*



The only Subscribers you need to include in a routing hierarchy are those that are used to pass Distributions on to other Subscribers. The end-node Subscribers (Subscribers that are only receiving and not passing on Distributions) do not need to be listed in the hierarchy. They have links in eDirectory to their parents.

Subscribers that exist in a routing hierarchy are generally parent Subscribers, although this is not required.

IMPORTANT: Subscribers that do not utilize parent Subscribers can increase the workload on the Distributor and increase network traffic across WAN links. All Subscribers should have a parent Subscriber, except for the first tier Subscribers that receive Distributions directly from the Distributor.

Rerouting Because of Changes to the Routing Hierarchy

If a parent Subscriber is changed, or the routing list (on the Routing Hierarchy tab of the Distributor object's properties) is changed, the change is reflected in the routing slip (data file used in the distribution process), because it is calculated each time the Channel schedule starts. A refresh is required for the Distributor to read eDirectory and obtain the new routing hierarchy.

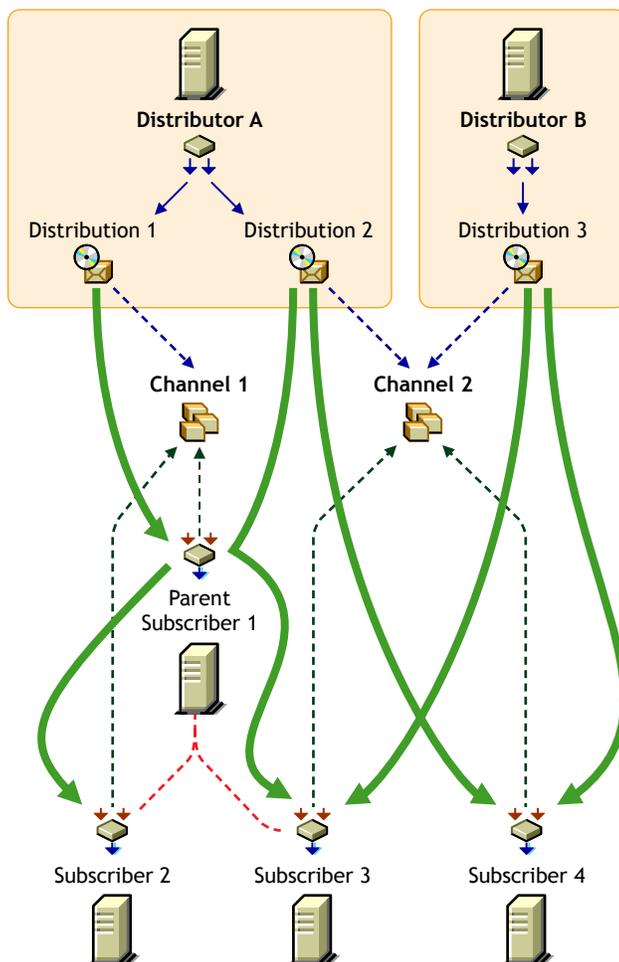
If a Subscriber server is removed from the network, and it was being used in a Distributor's routing hierarchy, you need to edit the Distributor object's properties to adjust the routing hierarchy because of that Subscriber's removal. Then refresh the Distributor so it can recognize the newer routing hierarchy.

Sharing Parent Subscribers with Other Distributors

If you have multiple Distributors, they can share portions of each other's distribution routes, meaning Subscribers can be listed in the distribution routing hierarchies of more than one Distributor. This is because the route to a Subscriber is dependent on the Distributor, and can be different for any given Distributor to Subscriber path.

Figure 3-10 illustrates the use of multiple Distributors and parent Subscribers in sending Distributions:

Figure 3-10 Using Multiple Distributors and Parent Subscribers



The arrows and lines indicate the subscription and Distribution connections to the Channels (dotted lines) and the distribution paths from the Distributors to the Subscribers (solid lines).

Figure 3-10 does not show distribution route hierarchies. For the purpose of this illustration, assume the following:

- ◆ Subscriber 1 is in Distributor A's hierarchy
- ◆ Subscriber 1 is a parent to Subscribers 2 and 3
- ◆ Subscribers 3 and 4 are in Distributor B's hierarchy
- ◆ Subscriber 4 is not in Distributor A's hierarchy

Note the following from the illustration concerning the use of multiple Distributors and parent Subscribers in sending Distributions:

- ◆ **Distribution ownership:** Distributors have ownership of their own Distributions and build and send each of their Distributions.
- ◆ **Multiple Distributors:** Multiple Distributors can list their Distributions in the same Channel. This means a Subscriber can receive Distributions from multiple Distributors.
- ◆ **Channel usage by Distributors:** Distributors can list their Distributions in any Channel, and they can list one Distribution in multiple Channels.
- ◆ **Multiple Distributions per Channel:** A Channel can have multiple Distributions from one or more Distributors.
- ◆ **Channel subscriptions:** Each Subscriber subscribes to any of the Channels that have the Distributions it needs. A Subscriber can subscribe to multiple Channels, and a Channel can have multiple Subscribers subscribed to it.
- ◆ **Parent Subscribers:** A parent Subscriber is used as a proxy for the Distributor to pass on Distributions to other Subscribers.
- ◆ **Orphaned Subscribers:** If a Subscriber is not in a Distributor's distribution route, or the child of a parent Subscriber in that hierarchy, the Distributor sends the Distribution directly to the Subscriber. This can be an issue for WAN links and other topology issues.

Distributing Across WAN Links

When you include parent Subscribers in the routing hierarchy, this can minimize network traffic by limiting the number of times a Distributor needs to pass a Distribution across a WAN link.

Because Distributors can send Distributions to parent Subscribers, which in turn can send them on to other Subscribers, a way is provided to send Distributions over a WAN link just once, instead of many times to reach every Subscriber on the other side of the WAN link.

Generally, you should have at least one parent Subscriber on every LAN to minimize the number of times a Distribution needs to cross a WAN link. Even if there are only two Subscribers on a LAN, you can reduce network traffic by using one of them as the parent Subscriber to the other.

Parent Subscribers are especially helpful with slow WAN links.

Consider the following when you determine how to distribute across your WAN links:

- ◆ **Parent Subscribers on the Distributor's LAN segment:** You should assign at least one Subscriber to be a parent Subscriber for all of the other Subscribers on a Distributor's LAN

segment. That way the Distributor can have more resources for sending Distributions across WAN links.

- ◆ **Parent Subscribers for bridging WAN links:** You can minimize the number of Subscribers that a Distributor must directly service across WAN links by assigning at least one parent Subscriber on all other LAN segments and including those parent Subscribers in the Distributor's routing hierarchy.

For example, your WAN has four LANs. With the Distributor in one LAN segment, it must send Distributions across three WAN links to get to Subscribers on the other three LAN segments. Let's assume each of the other LANs has 160 Subscribers that all need a Distribution from the Distributor. Without using parent Subscribers in the Distributor's routing hierarchy, the Distributor would need to send the Distribution 480 times across WAN links. In using parent Subscribers (four per LAN segment to share the Distribution workload on the LAN), the Distributor would only need to send the Distribution nine times.

- ◆ **Primary parent Subscribers on a LAN:** You can further minimize WAN traffic by tiering parent Subscribers on the other side of a WAN link from the Distributor. In other words, you can have just one parent Subscriber in the routing hierarchy that would also be a parent to several other parent Subscribers on its LAN segment.

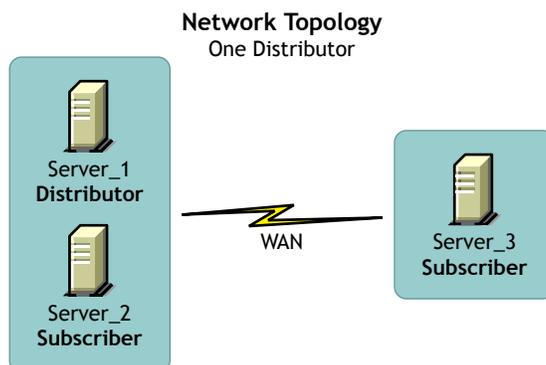
Using [Figure 3-10](#) as an example, Subscriber 1 on each LAN segment could be the parent Subscriber for Subscribers 2, 3 and 4. In turn, parent Subscribers 1, 2, 3, and 4 would each service their own Subscribers. That would allow the Distributor to just pass a Distribution across a WAN link once to Subscriber 1, which would take care of passing that Distribution on to the other three parent Subscribers, saving the Distributor three extra WAN link transmissions. Therefore, in contrast to [Figure 3-10](#), the 9 transmissions would be paired down to only three.

Out-of-Tree Distributions

To use Policy and Distribution Services in multiple trees, you must install the software separately in each tree. However, you only need to install the Server Management objects to one of the trees.

For example, if your network topology is as shown in [Figure 3-11](#):

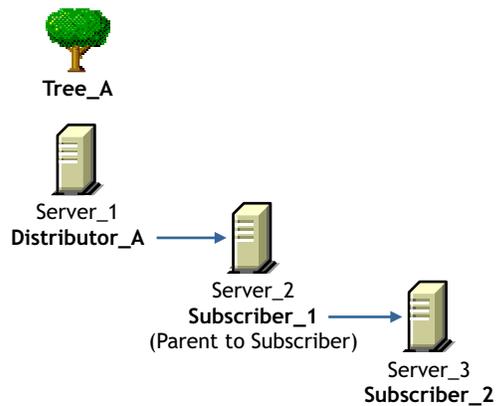
Figure 3-11 Network Topology with One Distributor



You could have the Tiered Electronic Distribution configuration shown in [Figure 3-12](#) for the Distributor's routing hierarchy:

Figure 3-12 *Distribution Flow in One Tree*

Distribution Flow
One Tree-One Distributor



In this example, the Subscriber and server objects all exist in Tree_A. This allows you to have centralized management of the Tiered Electronic Distribution objects, regardless of your network topology.

Although you can create the Distributor and Subscriber objects in only one tree, you can install the Policy and Distribution Services software to any server in your network, whether the server's eDirectory object resides in the same tree where the Tiered Electronic Distribution objects are created, or the server does not have an eDirectory server object in any tree (such as a Windows server in a domain). This allows you to have centralized management of Tiered Electronic Distribution in environments where you have multiple trees and mixed server operating systems (such as NetWare and Windows servers).

For information on how External Subscribers are used for sending Distributions between trees, see [Section 3.10.4, "Sending Distributions between Trees,"](#) on page 174.

Routing Hierarchy Configuration Guidelines

You should place parent Subscribers in the routing hierarchy using the following guidelines:

- ◆ Include at least one parent Subscriber on each LAN segment to minimize WAN traffic
- ◆ Include multiple parent Subscribers on each LAN that has 40 or more Subscribers to minimize a parent Subscriber's workload
- ◆ Make sure that every Subscriber that is not included in a Distributor's distribution route is assigned to a parent Subscriber on its LAN

Parent Subscribers are not always required for a WAN link. For example, if you have only two Subscribers on a LAN connected by a fast WAN link, the traffic difference between sending the Distribution once versus twice could be negligible. However, for a slow WAN link this might not be the case.

The factors in determining whether a Subscriber can receive Distributions directly from the Distributor instead of through a parent Subscriber are:

- ◆ Network connections

For example, are you distributing:

- ♦ only within a LAN?
- ♦ across a slow or fast WAN?
- ♦ across firewalls?
- ♦ in a NAT?
- ♦ Frequency of sending Distributions
- ♦ Size of the Distributions

3.3.3 Creating Distributors

By understanding your network's topology, your Distributions (how many, their sizes, and how often you might expect them to be rebuilt), and how many Subscribers receive the various Distributions, you can determine how many Distributors you need.

Distributors must be created by installing their software and eDirectory objects using the *ZENworks 7 Server Management with Support Pack 1 Program CD*. For more information, see “[Policy-Enabled Server Management Installation](#)” in the *Novell ZENworks 7 Server Management Installation Guide*.

To determine whether you need multiple Distributors, see [Section 1.1.4, “Are Additional Distributors Needed?”](#) on page 37.

3.3.4 Configuring Distributors

Distributor objects are automatically created when the Distributor's software is installed to a server. You can edit your Distributor object's properties at any time.

Not all properties associated with the Distributor object are required. Required properties are noted in the following steps; all others are optional.

- 1 In ConsoleOne, right-click the Distributor object, then click *Properties*.
- 2 Click *General > Settings* and fill in the following fields:

Use policy: Select to use the effective policy if you want to use the values set in the Tiered Electronic Distribution policy. This field is displayed if a Tiered Electronic Distribution policy has been created, distributed to the Distributor server, extracted by the Policy/Package Agent, and enforced on the server. If you select this option, the rest of the fields are dimmed and the policy settings are used instead.

Input rate: The rate Distributions received (for its Subscriber). The default is the maximum that the connection can handle. This rate is used to control a Distributor server's use of narrow bandwidth links.

Output rates based upon Distribution's priority: Sets the default output rate to minimize network traffic for Tiered Electronic Distribution objects. This determines the send rate for Distributors. The default value is the maximum that the connection can handle. Blank means that bandwidth is taken from third-party software.

There are three output priorities where you can specify a rate:

- ♦ **High priority:** These Distributions are sent before any Medium or Low priority Distributions.

- ♦ **Medium priority:** These Distributions are sent after all High priority and before any Low priority Distributions.
- ♦ **Low priority:** These Distributions are sent after all High and Medium priority Distributions.

For more information, see [Section 3.4.5, “Prioritizing Distributions,” on page 126.](#)

Maximum concurrent Distributions to build: Specifies the maximum number of distribution threads that can be running concurrently for building Distributions. The default value is 5. Valid values are from 1 to 10.

This number can help in load-balancing a Distributor’s building activity.

Maximum concurrent Distributions to send: Specifies the maximum number of distribution threads that can be running concurrently for sending Distributions. The default value is unlimited (a blank field).

This number can help in load-balancing a Distributor’s sending activity and spread network traffic over an entire scheduling window.

Connection time-out: Specifies the allotment of time before the Distributor server times out when connecting to another node. The default value is 300 seconds (five minutes), after which it ends the connection and does not retry until the send schedule starts again. The available range in seconds is 1 to 60,000.

You can increase or decrease this setting to allow messages to pass back and forth between the agents during the distribution process. If one node is expecting to receive a message from another, there should be a reasonable time to wait before assuming that the sender is no longer available.

IMPORTANT: This interval must be increased on slow or busy links where longer delays are frequent.

Working directory: Specifies the directory to be used by the Distribution. It contains Distributions, persistent status, and temporary working files. The default path is:

- ♦ **NetWare:** `sys:\zenworks\pds\ted\dist`

IMPORTANT: The default volume is `sys:` on NetWare servers. We recommend that you do not use the `sys:` volume because the directory’s content can become quite large.

- ♦ **Windows:** `c:\zenworks\pds\ted\dist`
- ♦ **Linux and Solaris:** `/var/opt/zenworks/zfs/pds/ted/dist`

The Distributor’s working directory is also used whenever a Distribution is created. A directory is created under the working directory using the DN of the Distribution object. For more information, see [Section 3.12, “Working Directories,” on page 186.](#)

3 Click *General > Messaging* and fill in the following fields:

Use policy: Select to use the effective policy if you want to use the values set in the Tiered Electronic Distribution policy. This field is displayed if a Tiered Electronic Distribution policy has been created, distributed to the Distributor server, extracted by the Policy/Package Agent, and enforced on the server. If you select this option, the rest of the fields are dimmed and the policy settings for messaging are used instead.

Server console: Specifies the level of output messages to send to the Distributor console on the server console.

For more information on the message notification levels, see [Section 3.11.5, “Minimizing Messaging Traffic,” on page 183](#).

SNMP trap: Specifies the level of messages to send via SNMP.

Log file: Specifies the level of messages to send to the log file.

Path and filename: You can specify a custom log file’s name and location for this Distributor object. The default is:

- ♦ **NetWare:** `sys:\zenworks\pds\ted\dist\ted.log`

IMPORTANT: The default volume is `sys:` on NetWare servers. We recommend that you do not use the `sys:` volume because the log file can become quite large.

- ♦ **Windows:** `c:\zenworks\pds\ted\dist\ted.log`
- ♦ **Linux and Solaris:** `/var/opt/zenworks/zfs/pds/ted/dist/ted.log`

For information on creating custom log files for all Distributor objects by using the Tiered Electronic Distribution policy, see [Section 4.3.5, “Creating Custom Log Files Using Policies,” on page 232](#).

Delete log entries older than __ days: Log file entries for a Distributor are deleted after they are older than the number of days specified. The default is six days.

E-mail: Specifies which level of messages are sent via e-mail.

Users: Specifies e-mail users for notification.

Address attribute: Specifies e-mail addresses for notification.

You can add users or groups stored in eDirectory or provide the e-mail addresses for users who are not contained in eDirectory. The e-mail Address Attribute associated with an eDirectory user is the default attribute.

IMPORTANT: If you select e-mail as a method for receiving notification, be aware that additional network traffic can be created.

4 Select the *Schedules* tab.

The schedule for a Distributor determines how often it reads the information contained in the Tiered Electronic Distribution objects in eDirectory. It reads the Channel, Distribution, and Distributor objects based on this schedule. You can set this up to reflect how often you expect information in these objects to change, or how often new objects might be created. However, we recommend that you leave its schedule set to the default of *Never* to prevent the possibility of the refresh interrupting the building or sending process, which could cause an infinite loop where the Distribution never finishes being built or sent.

We recommend that you after you create a Distribution, you force the Distributor to read eDirectory by right-clicking the Distributor object and selecting the *Refresh* menu option.

5 Select a schedule and fill in the fields:

Use policy: Select to use the effective policy if you want to use the values set in the Tiered Electronic Distribution policy. This field is displayed if a Tiered Electronic Distribution policy has been created, distributed to the Distributor server, extracted by the Policy/Package Agent, and enforced on the server. If you select this option, the rest of the fields are dimmed and the policy settings are used instead.

Schedule type: The Refresh schedule you selects determines when the Distributor reads eDirectory again.

IMPORTANT: We recommend the Distributor’s Refresh schedule be daily, unless changes to Distributions warrant a more frequent refresh. However, do not refresh the Distributor more often than every five minutes. The following can need up to five minutes to complete their processes: Distribution building, eDirectory replication, and tree walking (when no Search policy is defined).

For information on available schedules, see [Chapter 8, “Scheduling,” on page 317](#).

- 6 Select the *Routing* tab and create the Distributor’s routing hierarchy.

Subscriber routing hierarchy: Configure the routes the Distributor uses when sending Distributions to the Subscribers. You should plan this hierarchy in advance.

Use the following method to create the hierarchy:

- 6a Select the Distributor.
 - 6b Click *Add*, select one or more Subscribers, click *Select*, then click *OK*.
You can have multiple Subscribers directly under the Distributor.
 - 6c Select one Subscriber.
 - 6d Click *Add*, select one or more Subscribers, click *Select*, then click *OK*.
You can have multiple Subscribers directly under each Subscriber.
 - 6e Repeat [Step 6c](#) and [Step 6d](#) for each Subscriber until you have created the desired hierarchy.
- 7 Select the *Distributions* tab to view the Distributions being serviced by this Distributor.
 - 8 To edit a Distribution, select the Distribution, click *Details*, edit the properties, then click *OK* to exit the Distribution object’s properties.
 - 9 When you have finished configuring the Distributor and its Distributions, click *OK* to exit the Distributor object’s properties.

IMPORTANT: Changes made to Tiered Electronic Distribution objects (other than Distribution) are not in effect until the Distributor reads eDirectory.

3.3.5 Manually Refreshing the Distributor

Any time you make a change in eDirectory that affects the Distributor, you must manually refresh the Distributor so that it knows of that change.

For example, when you create a new Distribution, the Build schedule does not make the Distributor aware of the new Distribution. You must manually refresh the Distributor so that it can detect the change in eDirectory.

To refresh the Distributor:

- 1 In ConsoleOne, right-click the Distributor object.
- 2 Click *Refresh Distributor*.

This causes the Distributor to read eDirectory and obtain all of the changes that were made in eDirectory. The Distributor Agent can then act on any changes applicable to the Distributor.

To perform this task in iManager, see [“Forcing Policy and Distribution Services Agent Actions” on page 79](#).

Distribution building begins according to the current Build schedule. The Distribution is sent according to the Send schedule.

As soon as Subscribers receive an entire Distribution, they extract the contents to their working directories that are specified in the Subscriber objects' properties.

3.3.6 Deleting a Distributor Object and How Its Distributions Are Affected

You can delete Distributor objects from eDirectory. However, you can lose the following important information that you might want to reuse for the Distributor's replacement:

- ◆ The Distributor's distribution hierarchy
This is part of the Distribution object's properties, and it shows which Subscriber servers are used for passing on the Distributions.
- ◆ The list of its Distributions
The Distributor's Distributions become orphaned and unusable.

For information on how to handle orphaned Distributions, see [Section 3.4.10, "Handling Orphaned Distributions,"](#) on page 139.

3.4 Distributions

The following sections provide concepts and instructions for the Distribution object:

- ◆ [Section 3.4.1, "Understanding Distributions,"](#) on page 110
- ◆ [Section 3.4.2, "Distribution Issues,"](#) on page 114
- ◆ [Section 3.4.3, "Determining the Distributions,"](#) on page 117
- ◆ [Section 3.4.4, "Creating a Distribution,"](#) on page 123
- ◆ [Section 3.4.5, "Prioritizing Distributions,"](#) on page 126
- ◆ [Section 3.4.6, "Pre and Post Processing for Distributions,"](#) on page 126
- ◆ [Section 3.4.7, "Reassigning a Distribution to Another Distributor,"](#) on page 133
- ◆ [Section 3.4.8, "Deleting a Distribution,"](#) on page 136
- ◆ [Section 3.4.9, "Removing a Distribution Object - Auto Removal of Temporary Files,"](#) on page 137
- ◆ [Section 3.4.10, "Handling Orphaned Distributions,"](#) on page 139
- ◆ [Section 3.4.11, "Manually Importing and Exporting Distributions,"](#) on page 141
- ◆ [Section 3.4.12, "Using the Distribution Wizard,"](#) on page 143

3.4.1 Understanding Distributions

The Distribution (TED Distribution) object contains a list of data packages or data grouping information.

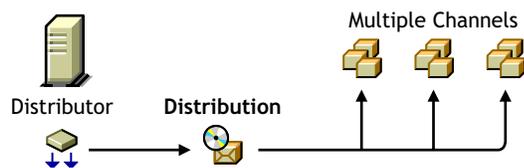
- ◆ ["Functional Relationship with Other Tiered Electronic Distribution Objects"](#) on page 111
- ◆ ["Distribution Description"](#) on page 111

- ◆ “Scheduling” on page 112
- ◆ “How New Versions of Existing Distributions are Created and Distributed” on page 112
- ◆ “Distribution Security” on page 112
- ◆ “Distribution Deletions” on page 113
- ◆ “Clean Up of Temporary Distribution Files” on page 113

Functional Relationship with Other Tiered Electronic Distribution Objects

Figure 3-13 illustrates a Distribution’s relationship with its Distributor and the Channels:

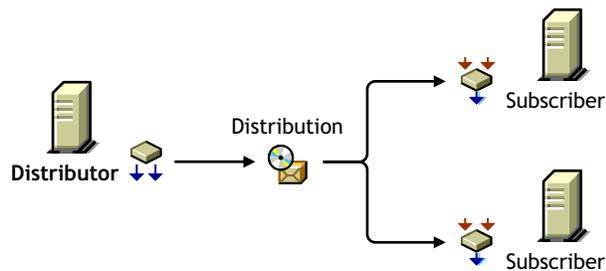
Figure 3-13 Distributor with a Distribution Listed in Multiple Channels



The Distributor associates its Distributions with the Channels.

Figure 3-14 illustrates that a Distributor sends Distributions to Subscriber servers:

Figure 3-14 A Distributor Sends Its Distribution to Two Subscribers.



Distribution Description

A Distribution is a compilation of software and/or files, or a policy package, that the various servers in your network might need.

A Distribution is owned by only one Distributor. A Distribution keeps a list of its Channel associations, and you can place the Distribution in multiple Channels.

When a Distribution is built, it is built according to its type. There are seven types of Distributions:

- Desktop Application ¹
- File
- FTP
- HTTP
- MSI
- Policy Package
- RPM
- Software Package

¹ The Desktop Application Distribution is only available when Desktop Management is installed.

For information on the different Distribution types, see [“The Distribution Types” on page 117](#).

Scheduling

A Distribution has a Build schedule that notifies its Distributor how often the Distribution needs to be built. If a Distribution has changed since the last time it was built, a new one is created.

Distributions can also be made active or inactive to control whether they should be built.

For information on scheduling, see [Chapter 8, “Scheduling,” on page 317](#).

How New Versions of Existing Distributions are Created and Distributed

After you have configured a Distribution object and set the various distribution schedules, newer versions of existing Distributions are automatically created and distributed according to the following parameters:

- ◆ **Refresh schedule:** This schedule determines when a Distributor reads eDirectory for changes to any of its Distributions. If changes are detected for a particular Distribution, it is rebuilt according to that Distribution’s Build schedule.

For more information on the Refresh schedule, see [“Distributor Object’s Refresh Schedule” on page 321](#).

- ◆ **Build schedule:** This schedule is set independently for each Distribution. When the schedule starts for a Distribution that has been determined to have had changes to it, the Distributor proceeds to rebuild that Distribution.

For more information on the Build schedule, see [“Distribution Object’s Build Schedule” on page 322](#).

- ◆ **Maximum revisions:** This field (in the Distribution object’s properties, click General > Settings), determines how many versions of a Distribution are kept on the Distributor and Subscriber servers’ file systems. For some Distribution types, this field determines whether a partial Distribution (delta) or complete Distribution is rebuilt. Otherwise, this field is used mainly to control disk space usage.

When the maximum number of revisions is being approached, an SMTP e-mail notification is sent, if SMTP notifications have been configured.

For more information on the Maximum Revisions field schedule, see [“Maximum Revisions” on page 115](#).

These parameters determine when a Distribution needs to be rebuilt. The other schedules (Send and Extract) determine when the rebuilt Distribution file is sent and extracted.

Distribution Security

Policy and Distribution Services provides several means for securing Distributions:

- ◆ [“Certificates” on page 113](#)
- ◆ [“Encryption” on page 113](#)
- ◆ [“Inter-Server Communications” on page 113](#)

Certificates

A certificate is a security mechanism used by Policy and Distribution Services to ensure that the Distribution received by a Subscriber was actually sent by the Distributor owning that Distribution. Without a matching certificate, a Subscriber cannot receive Distributions from the Distributor.

For more information, see [Section 7.1, “Distribution Security Using Signed Certificates and Digests,” on page 299](#).

Encryption

You can encrypt Distributions for when you send them outside your secure network.

For more information, see [Section 7.2, “Distribution Security Using Encryption,” on page 309](#).

Inter-Server Communications

You can secure communications between Tiered Electronic Distribution components residing inside and outside your secure network by installing inter-server communications security where needed.

For more information, see [Section 7.3, “Security for Inter-Server Communication Across Non-Secured Connections,” on page 313](#).

Distribution Deletions

When a Distribution is built, any deletions in the Distribution object or on the Distributor server’s file system, such as deleting files or directories, causes those files or directories to also be deleted from the Distribution when it is rebuilt. However, synchronization must be enabled in order for the files and folders to also be removed from the Subscriber server’s file system.

For more information, see [Section 3.11.1, “Directory Sync Granularity for File Distributions,” on page 176](#).

Clean Up of Temporary Distribution Files

To reduce the amount of disk space taken up by temporary Distribution files, those files are now automatically cleaned up. Previously, after a Subscriber server extracts a Distribution, the `distfile.ted` file is left in the Subscriber server’s working directory.

With the clean-up feature, you can have the `distfile.ted` file and other temporary Distribution files automatically deleted after the Distribution has been successfully extracted. However, so that the Subscriber is not sent the Distribution again, the status file is left in the Subscriber’s working directory to indicate that the Distribution has been extracted and cleaned up.

IMPORTANT: The Distribution clean-up feature works only for Distributions that have been both sent and received by ZENworks 7 Server Management Distributors and Subscribers. Distributions sent or received by Distributors or Subscribers running prior versions of Server Management software cannot be cleaned up.

- ◆ [“Parent Subscribers” on page 114](#)
- ◆ [“Distribution Types and Clean-Up” on page 114](#)
- ◆ [“Resetting Clean-Up Statuses in iManager” on page 114](#)
- ◆ [“Clean-Up Is Not Rollback” on page 114](#)

- ◆ “Clean Up of Reassigned Distributions” on page 114

Parent Subscribers

For parent Subscribers who might need to forward the Distribution, the files are not cleaned up on the parent Subscriber’s server, so that it can still forward the Distribution.

A parent Subscriber that has had a Distribution cleaned up which it is not forwarding, and then receives the same Distribution for forwarding, will receive the Distribution again, even though its status file indicates that it does not need it.

Distribution Types and Clean-Up

The FTP, HTTP, RPM, MSI, Software Package, and Policy Package types of Distributions are always cleaned up. The Maximum Revisions field is not available for those types of Distributions.

The File and Desktop Application types of Distributions can have their temporary files cleaned up after the Distribution has been extracted when the Maximum Revisions field is set to 1. However, for File Distributions, if the Verify Distributions check box is selected, the Distribution is not cleaned up, even if the Maximum Revisions field is set to 1.

Resetting Clean-Up Statuses in iManager

In iManager, you can reset the status file, which forces a Subscriber to receive a Distribution that has been cleaned it up.

Clean-Up Is Not Rollback

Cleaning up the temporary files does not cause any roll back of extracted Distributions. Clean-up is simply removing the temporary files.

Clean Up of Reassigned Distributions

The working directories for a Distribution that is reassigned from an old Distributor to a new Distributor are not automatically cleaned up on the old Distributor’s server. You need to manually clean up that Distribution’s temporary files on the old Distributor server.

3.4.2 Distribution Issues

Consider the following in determining your Distributions:

- ◆ File sizes and their potential for compression (.jpg files won’t benefit as much from compression as text files)
- ◆ The bandwidth of WAN links
- ◆ The frequency of file changes
- ◆ Network resource constraints, such as low disk space or extra bandwidth availability

The better you can determine this type of information, the better you can balance resource usage and minimize the use of resources.

You can configure Distributions to copy only files that are different than the target, or copy all files in their original state.

The following sections provide information about other issues with Distributions:

- ♦ [“Maximum Number of Concurrent Distributions” on page 115](#)
- ♦ [“Maximum Revisions” on page 115](#)
- ♦ [“I/O Rate \(Bytes per Second\)” on page 116](#)
- ♦ [“Updating the Distributor’s eDirectory Information” on page 116](#)
- ♦ [“Checking the Distribution Package Changes” on page 116](#)
- ♦ [“MSI Distribution Extraction Errors” on page 116](#)

Maximum Number of Concurrent Distributions

This is an attribute found in the Distributor and Subscriber objects. It is used to control the number of Subscribers that can be serviced concurrently when sending Distributions. This is helpful if the Distributor or parent Subscriber is servicing a large number of Subscribers. It prevents the Distributor from spreading itself very thin and sending the Distribution to all of the Subscribers at once.

For example, if a Distributor or parent Subscriber sends to 100 Subscribers and the number of concurrent Distributions is set to 10, then the sender starts with 10 connections. As one connected Subscriber finishes receiving the Distribution, another Subscriber is added in its place in the list of 10. This continues until all 100 have been serviced.

Maximum Revisions

Each Distribution allows you to determine how many versions of the Distribution are kept by the Distributor and Subscribers in their working directories. The default is infinite for all Distribution types, whether the Distribution is created in ConsoleOne or iManager; however, the File and Desktop Application types of Distributions have a default of 10 if they are created in ConsoleOne. Make sure that you fill in the Maximum Revisions field attribute when creating Distributions. Consider disk space availability when calculating the maximum number of revisions.

If you select to limit the revisions, the e-mail fields are available, where you can specify a trigger to notify e-mail recipients when your maximum number is approached, as well as define who the e-mail recipients are. If you select to allow unlimited revisions, the e-mail fields are not available.

The File Distribution only builds a complete Distribution the first time it creates the Distribution. All subsequent versions are just the differences (deltas) between a current version and its previous version. However, when the File Distribution reaches its maximum number of revisions, it deletes all previous versions and build an entirely new Distribution (called a baseline), and starts from 1 in counting the number of revisions.

When the maximum number of revisions is met for FTP, HTTP, and Server Software Package Distribution types, the agent deletes the oldest version of the Distribution and adds the current version to the revisions. Therefore, it never exceeds the maximum number entered in the Distribution object.

When the maximum number of revisions is being approached, an SMTP e-mail notification is sent if SMTP notification has been configured.

I/O Rate (Bytes per Second)

This is an attribute found in the Distributor and Subscriber objects. It is used to control the amount of bandwidth used by the Distributor or parent Subscriber when sending Distributions. The default is unlimited, meaning the sender uses all the bandwidth available in sending Distributions.

Updating the Distributor's eDirectory Information

The Distributor must be updated with the configuration information contained in the Tiered Electronic Distribution objects in eDirectory.

Configuration changes include any changes made to the attributes of the Distributor object, Distribution objects belonging to that Distributor object, or Channel objects to which the Distributor object is associated.

The Distributor has a schedule that determines how often it reads eDirectory for configuration information. Set this schedule to coincide with the frequency at which Tiered Electronic Distribution objects are modified in eDirectory.

You can also force an eDirectory refresh by right-clicking a Distributor object and selecting the Refresh menu option, or by using the ZENworks Server Management role in iManager (see [“Forcing Policy and Distribution Services Agent Actions” on page 79](#)).

Checking the Distribution Package Changes

The Distribution's Build schedule tells the Distributor the frequency at which the Distribution should be checked for changes.

For example, the Distribution schedule might specify a weekly build. The Distributor rebuilds that package and compares it to the previous version to see if there have been any changes.

MSI Distribution Extraction Errors

Some MSI Distributions can fail to extract on Windows 2000 servers (usually displaying error 1603), but not on Windows Server 2003 servers. The difference is in how the two operating systems differently handle the rights needed to install the MSI packages.

This can be solved for Windows 2000 servers by editing the properties in the MSI Distribution:

- 1 In ConsoleOne, access the MSI Distribution's properties.
- 2 On the *Type* tab, select the MSI package listed under *Selected Packages*.
- 3 Click *Edit Parameter List* to open the Edit Parameters dialog box.
- 4 In the *Custom Parameters* field, enter:
`ALLUSERS=1`
- 5 Click *OK* to save the change.
- 6 Repeat [Step 2](#) through [Step 5](#) for each MSI package listed under *Selected Packages*.
- 7 Click *OK* to save the updated MSI Distribution properties.

3.4.3 Determining the Distributions

You can distribute whatever you can represent on the file system. This includes server applications and files. For example, the applications or files could fulfill one of the following purposes:

- ◆ Installing server software (such as virus protection software)
- ◆ Updating server software (such as a NetWare support pack)
- ◆ Updating files (such as virus patterns) on servers
- ◆ Enforcing standardization of server files or configurations (such as replacing the `autoexec.ncf` file on a NetWare server with an updated version)

Use a descriptive method for naming the Distributions. You can also use these names for naming the Channels associated with the Distributions. For example:

```
VirusProtect
VProtectPatterns
NW51patch4
NW6patch1
AUTOEXECNCF000326
```

The following sections explain the different Distribution types and issues related to determining your Distributions:

- ◆ [“The Distribution Types” on page 117](#)
- ◆ [“Determining the Sizes and Frequencies for Distribution Packages” on page 122](#)

The Distribution Types

There are several Distribution types. Each type has unique features that tailor it for specific needs.

- ◆ [“Desktop Application” on page 117](#)
- ◆ [“File” on page 118](#)
- ◆ [“FTP” on page 119](#)
- ◆ [“HTTP” on page 120](#)
- ◆ [“MSI” on page 120](#)
- ◆ [“Policy Package” on page 121](#)
- ◆ [“RPM” on page 121](#)
- ◆ [“Software Package” on page 122](#)

For information on how to configure each Distribution type, see [Section 3.4.4, “Creating a Distribution,” on page 123](#) (specifically, [Step 6 on page 124](#)).

For the File and FTP types of Distributions, a Distribution Wizard is available for automating the process of creating them. For more information, see [Section 3.4.12, “Using the Distribution Wizard,” on page 143](#).

Desktop Application

Distributes the Application objects (that are created in Desktop Management) and the application’s associated files to specified locations on the eDirectory tree and target Subscriber servers. This Distribution type allows you to solve geographic, workload, and redundancy issues for applications

distributed by Novell Application Launcher that otherwise might require much of your time in manual configuration work in Desktop Management. For more information, see [Chapter 6, “Desktop Application Distribution,” on page 273](#).

The Desktop Application Distribution type is not supported for Linux and Solaris servers.

This Distribution type automatically distributes a modified copy of the original Application object to a context in the eDirectory tree (a Subscriber’s working context), and automatically copies the application’s files to the Subscriber server that can locally service its users and workstations. It performs all of the appropriate hookups to the modified Application object to render it fully functional.

For the Desktop Application Distribution, you can set the maximum number of revisions in the Distribution object. When the version number reaches the number that you set, the Distributor rebuilds the entire Distribution. By default, this number is 10.

You can send Desktop Application Distributions to Subscriber servers on a tree that is different from the Distributor server’s. However, the recipient server’s Subscriber object must reside in the same tree as the modified Application objects that are created by the Distribution. The External Subscriber object is used on the Distributor’s tree to send a Desktop Application Distribution to a server on another tree.

File

With this type you can select files and/or directories from the Distributor server’s file system for distribution, and select a destination location for extraction on the Subscriber.

The File type is sequential, meaning it controls the order for the building and extraction of Distributions. This prevents the building and extracting processes from being performed out of sync.

IMPORTANT: Linux and Solaris file systems are case sensitive to allow paths and filenames that are identical except for case differences. However, if you select two such files, only the first file selected during extraction is distributed, because the File type is not case sensitive. Therefore, do not place two files into a File Distribution where their paths and filenames are identical except for case differences.

Also, if a NetWare server is the target for a File Distribution, you might encounter an error due to code page differences where extended characters are used (such as ê, ë, ì, or í). The information in [“Extended Characters in Directory Paths” on page 287](#) in the [Desktop Application Distribution](#) section is also applicable to File Distributions.

By default, Cache and Forward is used. This process allows a parent Subscriber to begin sending a Distribution to subordinate Subscribers before it has finished receiving the Distribution. This allows entire Distributions to be sent more quickly through a chain of parent Subscribers in the Distributor’s routing hierarchy than if they each had to wait until each Subscriber had completed receiving the Distribution before it started sending.

The File Distribution is useful for distributing large Distributions that change often, thus requiring updates that need to be distributed frequently.

For the first version of a Distribution, the Distributor builds the entire Distribution (creating a baseline). A unique feature of the File type is that for all subsequent versions it calculates the differences at build time and only builds a delta of the Distribution.

The File type does this by keeping a list of the files and directories contained in a Distribution on the source machine (the Distributor or a parent Subscriber). If a source file changes, a new Distribution is built the next time its Build schedule starts. However, this new Distribution only contains the files that are different between the previous version and the current version. This is known as a delta of the original Distribution.

This delta of the Distribution file is what is distributed to the Subscribers-not the entire Distribution.

The File type is also effective when changes are frequent because it can build much smaller deltas.

There is no option to send the entire File Distribution. However, after the maximum number of revisions has been met, the Distribution is completely rebuilt and all deltas and previous revisions are deleted. Therefore, if you set the maximum number of revisions to 1, deltas are not used and the entire Distribution is built and sent every time.

For example, the first build is the baseline Distribution (version 1), the first update (Delta 1) is version number 2, the second update (Delta 2) is version number 3, and so on until the number of revisions you set is reached, which triggers a new baseline rebuild. By default, this number is 10.

Pre and Post actions can be set for File Distributions. For more information, see [Section 3.4.6, “Pre and Post Processing for Distributions,” on page 126](#).

You can set the maximum number of revisions in the Distribution object.

If synchronization is enabled, you can use the File type for removing files and directories from the Subscriber server’s file system upon extraction of the Distribution in one of two ways:

- ◆ **Edit the Distribution object:** Remove files from the list of files and directories in the Distribution object. When the Distribution is built again, those files and directories are not included.
- ◆ **Remove files from the Distributor’s file system:** Remove files from the Distributor’s file system that were part of the Distribution. When the Distributor is refreshed, it rebuilds the Distribution without those files and directories.

In both cases, upon extraction of the Distribution, and with synchronization enabled, those files and directories are removed from the Subscriber server’s file system. For more information on synchronization, see [Section 3.11.1, “Directory Sync Granularity for File Distributions,” on page 176](#).

To manually force a Distribution to be built, you can use iManager (see [“Forcing Policy and Distribution Services Agent Actions” on page 79](#)).

FTP

With this type you can create a Distribution consisting of files from one or more FTP sources. Each source can contain one or more directories and/or files.

When an FTP site directory entry is a directory, all of its files and subdirectories are built for the Distribution.

Server Management now supports retrieval of symbolic link files. This allows the Linux or Solaris environments to receive FTP files that would be considered invalid on other platforms.

Whenever a Distribution's Build schedule starts:

- ◆ The FTP type creates a new Distribution only if the new version would be different than the previous version.
- ◆ The Distributor builds the entire new Distribution.
- ◆ The Distributor sends each new version of the Distribution to the appropriate Subscribers.

You can set a maximum number of revisions in the Distribution object to conserve disk space. By default, the number is unlimited.

HTTP

With this type you can create a Distribution consisting of one or more HTTP sources. Each source can contain one or more target entries.

Whenever a Distribution's Build schedule starts:

- ◆ The HTTP type creates a new Distribution only if the new version would be different than the previous version.
- ◆ The Distributor builds the entire new Distribution.
- ◆ The Distributor sends each new version of the Distribution to the appropriate Subscribers.

You can set a maximum number of revisions in the Distribution object to conserve disk space. By default, the number is unlimited.

MSI

Distributes Microsoft Software Installer (MSI) packages to Windows servers, where the MSI engine is used to install the Windows-specific software included in an MSI Distribution. Any vendor can create MSI packages for their software for installing in a Windows environment.

The Installshield* AdminStudio* ZENworks Edition software for creating .mst files is available on its own CD that is provided with the ZENworks 7 product.

MSI 3 is supported as a version that can be distributed. However, ZENworks 7 does not individually support any of MSI 3's new features.

The components of an MSI Distribution consist of .msi, .msp, and .mst files:

- ◆ **.msi file:** An MSI package containing Microsoft software to be installed by the MSI engine.
An MSI package can include just the .msi file, or the .msi file with the other files in its folder and all of the files contained in any subfolders.
- ◆ **.mst file:** A file that adds, deletes, or changes the properties in an MSI package to enable customizing of the installation for different groups of users.
- ◆ **.msp file:** An MSP package that provides a patch to an MSI package.
MSI-based patch files might have filename extensions other than .msp.

An MSI Distribution might contain:

- ◆ One or more MSI packages
- ◆ One or more MSI packages with one or more .mst files applied to each MSI package
- ◆ One or more MSI packages with one or more MSP packages

- ◆ One or more MSI packages with one or more MSP packages and one or more .mst files applied to each MSI package
- ◆ One or more MSP packages only, because .msp files contain the information necessary for identifying the MSI packages' applications that they are to patch

Patching can include modifying the settings of a machine, as well as updating files.

You can determine the application order of the .mst files for each MSI package, and you can determine the execution order of the MSI and MSP packages listed in the Distribution.

When an MSI Distribution includes both of the MSI and MSP components, post-installation actions are added by the Distributor to the Distribution to ensure the correct order of completion.

Because an MSP is designed to modify a specific MSI package, you need to make sure that you have the correct order of execution.

Some MSI Distributions can fail to extract on Windows 2000 servers. To solve this problem, see [“MSI Distribution Extraction Errors” on page 116](#).

IMPORTANT: Because an MSI Distribution recursively gathers all of the files from the MSI file's location, if you have multiple .msi files in a given location, all other files and subdirectories contained therein are gathered once for each .msi file. The distribution gathering process cannot determine which other files or subdirectories belong to each .msi file, so you can end up with a much larger MSI Distribution file than is necessary. Therefore, instead of storing your .msi files in one directory, place each into its own subdirectory with its own supporting files and subdirectories.

Policy Package

This type provides the mechanism for applying policies to servers. In previous versions of Policy and Distribution Services, all policies were enforced through eDirectory object and container associations. With ZENworks 7 Server Management, policies are now distributed Subscriber servers for enforcement using the Distributed Policy Package. However, policies for Distributors continue to be enforced through context associations using the Container Package or Service Location Package.

With the Policy Package Distribution, you send policies directly to servers as Distributions, which are extracted on the receiving Subscriber server. The contained policies are then enforced on that server.

You can set a maximum number of revisions in the Distribution object to conserve disk space. By default, the number is unlimited.

To send a Policy Package Distribution to a Subscriber using an External Subscriber object, you must edit the `agentinfo.properties` file to prevent trusted tree errors. For more information, see [“Preventing Trusted Tree Errors for Policy Package Distributions” on page 161](#).

For more information on each policy, see [Section 4.1.6, “Server Policy Descriptions,” on page 200](#).

RPM

You can distribute any Red Hat Package Manager (RPM) packages that you have previously created to your Linux and Solaris servers using the RPM Distribution.

For Solaris, RPM must first be installed on the server, because it is not installed with Solaris software by default. Solaris' equivalent is PKG.

Whenever a Distribution's Build schedule starts:

- ◆ The Distributor builds the entire new Distribution.
- ◆ The Distributor sends each new version of the Distribution to the appropriate Subscribers.

You can set a maximum number of revisions in the Distribution object to conserve disk space. By default, the number is unlimited.

Software Package

A Server Software Package is created in ConsoleOne in the Server Software Package namespace. For more information, see [Chapter 5, "Server Software Packages," on page 237](#).

Software Package is the most robust type of Distribution. It includes installation prerequisites, pre-installation instructions, post-installation instructions, and the ability to modify text fields, SET parameters, registry settings, and the `products.dat` file.

With the Software Package Distribution you can select `.cpk` files for distribution. This allows you to place a software product into a Distribution for automatic installation on the receiving server. This can include software updates to existing server software on the server.

You can select multiple `.cpk` files for one Distribution. Then, individual `.cpk` files are applied on the Subscriber, depending on whether the `.cpk` file's prerequisites are met.

IMPORTANT: The order that the `.cpk` files are applied on a server is not guaranteed, and `.cpk` files contained in one Distribution that might start in a certain order might not all finish in that same order. Therefore, place each `.cpk` file in its own Distribution if you want them to be installed in a particular order and use Distribution scheduling to determine the order. For more information, see ["Forcing the Software Package Distribution Order Using Multiple Distributions" on page 239](#).

Determining the Sizes and Frequencies for Distribution Packages

A Distribution's size and frequency of being built and sent depends on the following:

- ◆ The size and number of files being distributed. Knowing this helps in determining the amount of disk space to be used on Distributor, Subscriber, and parent Subscriber servers.
- ◆ A Software Package Distribution (`.cpk`) always builds an entirely new version of the Distribution each time the source changes.
- ◆ HTTP and FTP Distributions always build an entirely new version of the Distribution whether the source has changed or not.
- ◆ How often the packages change and need updating. Knowing this helps determine how frequently new versions of the package are created. Servers required to rebuild large Distribution packages on a regular basis should have the processing power to perform this work. The creation of many versions of a package also affects the amount of disk space used in the Distributor's working directory.
- ◆ The number of versions of a Distribution package that are retained. This also affects disk space usage on the Distributor's and Subscribers' servers.
- ◆ The File Distribution creates a delta file for each new version of the Distribution until it reaches the number you have specified in the Maximum Number of Revisions field (10 is the default). Then it begins a new baseline Distribution. The delta file contains only the differences between the last and current versions of the Distribution.

3.4.4 Creating a Distribution

1 In ConsoleOne, select the container where you want the Distribution to be created, click *File > New > Object*, select the *Distribution* type, then click *OK*.

2 Specify a Distribution name.

IMPORTANT: Periods (.) are not allowed in Distribution names. Instead, use dashes (-) or underscores (_) as word separators. If you use a period in the Distribution name, the Distribution is not sent, and the Distributor is not reloaded after it has been exited.

3 To give the Distributor ownership of the Distribution, browse to select the Distributor object, select *Define Additional Properties*, then click *OK*.

The Distribution object's properties are displayed.

Each Distribution belongs to a single Distributor that builds and sends the Distribution.

4 Click *General > Settings* and fill in the fields:

Active: Required. In order to make a Distribution available to Subscribers, it needs to be active.

Use digests: Digests are used by Distributors and Subscribers to verify that Distributions have not been tampered with while in transit. The digest provides an MD5 checksum for the Subscriber to compare.

Creating a digest takes more time on larger Distributions. The number of minutes per megabyte is dependent on the hardware configuration of the server where the digest is being created.

Digests also detect corruption in a Distribution's package. In the case of corruption, the Subscriber renames the `distfile.ted` Distribution file to `distfile.corrupt` and the Distribution is rebuilt and sent the next time the Channel's schedule fires.

Encrypt: You can have the Distribution encrypted if you are sending it across non-secured connections. Encryption provides security for the Distribution during transit between the Distributor and Subscriber when they are not within the same firewall. Select either Strong or Weak encryption. You also must have the same version of NICI 2.6.4 installed to each of these servers for encryption to work (see ["Installing NICI 2.6.4" on page 54](#)). However, if you already have NICI 2.4.6 installed, it is optional whether you upgrade to NICI 2.6.4, because these versions are compatible with each other.

Maximum revisions: This number helps you to control disk space usage by determining how many versions of a particular Distribution are kept in the Distributors' and Subscribers' working directories. The default is 10. Select *Limited* and enter a number.

Increase the number if data is changing often and the changes are minimal (smaller delta files). Decrease the number if data is not changing very often, or if a significant amount of data is changing (larger delta files).

The following e-mail options are available if you set a maximum number. If you select *Unlimited*, these options are dimmed.

- ♦ **Approaching maximum revision email notification list:** Contains the e-mail addresses of anyone who is to be notified when a Distribution is approaching the maximum revisions set in the *Maximum revisions* field. Here, you can either remove a single or all displayed addresses.
- ♦ **Email address (maximum revision notification):** You can add e-mail addresses to the list in *Approaching maximum revision email notification list*. Just enter an e-mail address and click *Add* and it is displayed in the listing.

- ◆ **Send notifications when Distribution revision is ___ or less of reaching maximum revisions:** Enter a number to indicate how close “approaching” is. When the current revision number of Distribution plus this number equal the maximum revisions, an SMTP notification is sent to the listed addressees.

SMTP must be configured and its e-mail server address listed in the next field.

- ◆ **Email server address:** The SMTP server used to send the e-mail notifications. For example, mail.novell.com.

For information on configuring SMTP e-mail notifications, see [“SMTP Host” on page 208](#).

Priority: You can give the Distribution a priority that determines how it is sent in relation to other Distributions. A High priority means it is sent before Medium or Low priority Distributions. For information on prioritizing Distributions, see [Section 3.4.5, “Prioritizing Distributions,” on page 126](#).

Distributor: The DN of the Distributor object that builds and sends this Distribution. This attribute cannot be modified. You selected the Distributor when you created the Distribution object.

Description: Provide useful details about the Distribution, such as the name of the desktop application, the files and directories it contains, intended user groups, and so on.

- 5 Click *General > Restrictions* and select a platform restriction:

Platform restrictions: If you want to select specific operating system versions as a prerequisite to receiving this Distribution, deselect No Restrictions and select the desired operating system version. You can select from the following:

- No Restrictions
- NetWare All
- NetWare 4.x (earlier versions of ZfS supported these platforms)
- NetWare 5.0 (earlier versions of ZfS supported this platform)
- NetWare 5.1
- NetWare 5.x
- NetWare 6.x
- Windows Server
- Solaris
- Linux

Selecting the No Restrictions check box means that the Distribution can be sent to any platform.

If you select NetWare All, you do not need to select any of the individual NetWare platforms.

- 6 Select the *Type* tab and use the drop-down box to choose a Distribution type in the *Select Type* field:

[Section A.1, “Desktop Application,” on page 383](#)

[Section A.2, “File,” on page 383](#)

[Section A.3, “FTP,” on page 387](#)

[Section A.4, “HTTP,” on page 390](#)

[Section A.5, “MSI,” on page 392](#)

[Section A.6, “Policy Package,” on page 395](#)

[Section A.7, “RPM,” on page 396](#)

[Section A.8, “Software Package,” on page 397](#)

For some Distribution types, when entering information into a field, such as a directory name, be sure to press Enter or the change is not saved.

IMPORTANT: For the FTP, HTTP, RPM, Software Package, and Desktop Application types of Distributions, if a target file is found to be locked during extraction, the Subscriber throws an exception stating that the file could not be copied. The Distributor receives this information from the Subscriber and logs the failure in the reporting database.

7 Select the *Schedule* tab and select a schedule:

The Build schedule determines how often the Distributor builds a new version of the Distribution.

Two options allow you to override the Channel's Send and Subscriber's Extract schedules:

- ♦ **Send Distribution immediately after building:** Overrides the Channel's Send schedule, allowing you to immediately send the Distribution, rather than wait for the Send schedule to start. However, the Subscriber's Extract schedule determines when it is extracted for use.
- ♦ **Extract Distribution immediately after receiving:** Overrides the Subscriber's Extract schedule, allowing the Distribution to be immediately extracted, rather than wait for the Extract schedule to start. This is useful for Distributions that need to be extracted immediately, such as a Distribution that provides virus patterns.

Build schedule for File Distributions: This type builds a new Distribution and compares it with the previous version for changes. If there are changes, the File type builds a file consisting of the differences between the current version and the previous version. When the maximum number of versions is reached, the type builds a complete Distribution (not just a file containing the differences) and deletes all previous versions.

Build schedule for HTTP, FTP, and Software Package Distributions: These types build new versions of the Distribution each time the Build schedule starts, regardless of whether the Distribution has changed. It sends this new version to all Subscribers.

When sending a Distribution, the sender retries every 2 minutes for 30 minutes, then stops. It does not begin sending again until the Channel schedule starts again.

8 Select the *Channels* tab and fill in the field:

Channels: Each Distribution must be associated with at least one Channel in order for it to be sent to a Subscriber. A Distribution is sent to all Subscribers of the selected Channel or Channels.

9 If you want to set pre or post actions for the Distribution, see [Section 3.4.6, "Pre and Post Processing for Distributions,"](#) on page 126 for the steps.

10 Click *OK*, then select *Yes* to resolve the certificates.

For NetWare and Windows servers, this copies the security certificates from the Distributor to Subscriber subscribed to the Channel. For Linux and Solaris servers (if you do not have drives mapped to them), you may need to resolve the certificates manually.

For information, see [Section 7.1.6, "Resolving Certificates,"](#) on page 303.

3.4.5 Prioritizing Distributions

Distributions can be prioritized in two ways:

- ◆ **Send queue:** You can prioritize the order in which Distributions are sent: High, Medium, or Low. For example, in a given Channel, all High priority Distributions are sent first, then the Medium priority Distributions are sent, and then the Low priority Distributions are sent.

Because Distributions with mixed priorities cannot be sent concurrently, you can control the order in which Distributions are sent by the priorities that you assign them.

- ◆ **Output rate:** You can configure different output rate settings for a Distribution, based on a priority: High, Medium, or Low. This allows you to control the bandwidth a Distribution uses. For example, if you want your High priority Distributions to utilize the most bandwidth, you should configure their output rates with the High priority. Blank means that bandwidth is taken from third-party applications.

The Maximum Number of Concurrent Distributions value is affected by prioritizing. This value is subordinate to the priorities set for the Distributions. For example:

- ◆ You have the concurrent Distribution number set to 10.
- ◆ There are 3 High priority Distributions.
- ◆ There are 6 Medium priority Distributions.
- ◆ There are 20 Low priority Distributions.
- ◆ Initially, only the 3 High priority Distributions are sent concurrently.
- ◆ After all 3 of the High priority Distributions are sent, the 6 Medium priority Distributions are sent concurrently.
- ◆ After all 6 of the Medium priority Distributions are sent, 10 of the 20 Low priority Distributions are sent concurrently, and so on.

3.4.6 Pre and Post Processing for Distributions

Pre and post processing actions are new features for Distributions in ZENworks 7 Server Management:

- ◆ [“Pre and Post Processing Actions Now Available in Distributions” on page 126](#)
- ◆ [“Pre and Post Actions in Software Packages versus Distributions” on page 127](#)
- ◆ [“The Pre and Post Feature Enhances Software Package Distribution Processing” on page 127](#)
- ◆ [“Error Messages Given When Valid Distribution Types Are Not Selected” on page 128](#)
- ◆ [“Pre and Post Distribution Processing Actions” on page 128](#)

Pre and Post Processing Actions Now Available in Distributions

To apply execution logic to a Distribution, pre and post actions are now available for the following Distribution types:

File
FTP
HTTP
MSI

RPM

Software Package 1

1 Previously, only a Server Software Package had this functionality. Now both the software package and its Software Package Distribution can have pre and post actions defined.

The benefit of having pre and post actions in these Distribution types is that you are no longer restricted to using only Server Software Packages to perform those actions.

The pre and post processing actions are not available for the following Distribution types:

Desktop Application

Policy Package

Pre and Post Actions in Software Packages versus Distributions

In Server Software Packages, the pre and post features are contained in two different tabs: Pre-Installation and Post-Installation, with Script and Load/Unload tabs for accessing the various options.

For the Distribution types that now have this feature, a Pre/Post Actions tab has been added to their Distribution object's properties (with Pre-Distribution Actions and Post-Distribution Actions tabs).

The following options are available from the Pre/Post Actions tab:

Load Java Class

Script

Start Process

Stop Process

Start Windows Service

Stop Windows Service

For more information on these options, see [“Pre and Post Distribution Processing Actions” on page 128](#).

The Pre and Post Feature Enhances Software Package Distribution Processing

When either a pre or post action is defined for a Software Package Distribution, the following is done:

1. A list of .cpk files contained in the Distribution is created in the Type tab of the Distribution object.
2. All pre actions are processed according to the order you defined for them.
3. The .cpk files are processed serially.
4. All post actions are processed according to the order you defined for them.

You can use Pre and Post Actions in a Distribution object containing multiple software packages to ensure pre and post actions are performed before and after the software packages listed in the Distribution are processed. However, pre and post processing only guarantees the order on ZENworks 7 Server Management Subscribers, because this functionality is not backwards compatible with ZENworks for Servers 3.x Subscribers.

Error Messages Given When Valid Distribution Types Are Not Selected

There are some instances when the Pre/Post Actions tab display only a message:

- ◆ In a Distribution object's properties, if you have not yet selected a Distribution type, the following message is displayed on the Pre/Post Actions tab:
You must select a Distribution type before you can configure pre or post actions.

However, you must not only select a Distribution type, you must also "save" it by clicking Apply. Then the Pre or Post page recognizes the Distribution and the Pre or Post actions can be applied.
- ◆ If the Distribution type you have selected is either Policy Package or Desktop Application, the following message is displayed on the Pre/Post Actions tab, because pre/post actions are not supported for those types of Distributions:
This Distribution type does not support pre or post distribution actions.

Pre and Post Distribution Processing Actions

In each of the following sections, the information provided applies to both the Pre-Distribution Actions and Post-Distribution Actions subtabs of the Pre/Post Actions tab. The difference is that Pre-Distribution Actions occur before the main Distribution is extracted and Post-Distribution Actions occur after the Distribution has completed extracting.

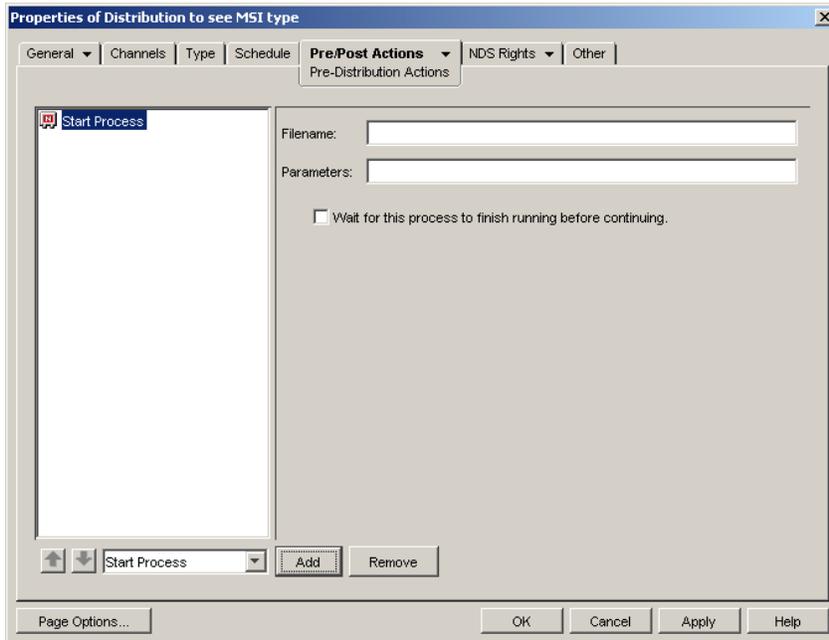
- ◆ "Start Process Action" on page 128
- ◆ "Stop Process Action" on page 129
- ◆ "Start/Stop Windows Service Action" on page 130
- ◆ "Script Action" on page 132
- ◆ "Load Java Class Action" on page 133

Start Process Action

This action works for Windows services, Java processes, and NLM processes.

The Start Process action is similar to the Load NLM/Process action in Server Software Packages, as illustrated in [Figure 3-15](#):

Figure 3-15 Properties of Distribution to See MSI Type Dialog Box



To add a Start Process action, select the option in the drop-down box and click the Add button. Then fill in the fields:

- ◆ **Filename:** This must be the exact name. For NetWare, include the .nlm extension.
For Linux and Solaris, you must include the full path.
- ◆ **Parameters:** Include any command line parameters for the NLM™ or process being run.
- ◆ **Wait for this process to finish running before continuing:** You can select this option for an NLM or process that terminates itself. It must terminate within 10 minutes, or the whole loading process fails. By default, this option is deselected.

If you select an NLM to be loaded by the Distribution, and the NLM is already running on the target server, the package installation fails and is rolled back (if rollback is enabled).

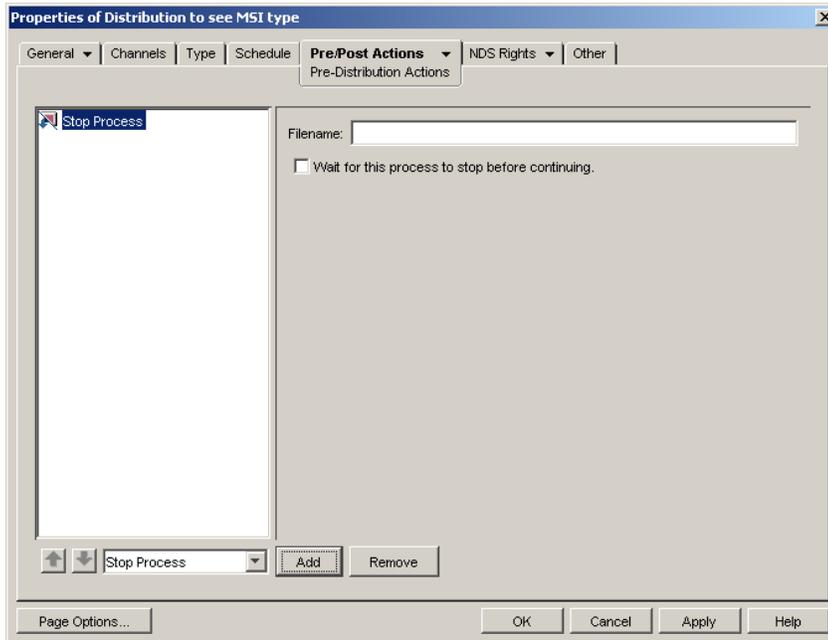
You can make sure that an NLM is not already loaded when you are including it in the Distribution by adding a Stop Process option for that NLM before adding the Start Process option-but only if this NLM does not require user input from the keyboard to unload it.

Stop Process Action

This action works for Windows services, Java processes, and NLM processes.

The Stop Process action is similar to the Unload NLM action in Server Software Packages, as illustrated in [Figure 3-16](#):

Figure 3-16 Properties of Distribution to See MSI Type Dialog Box



To add a Stop Process action, select the option in the drop-down box and click the Add button. Then fill in the fields:

- ◆ **Filename:** This must be the exact name, including the extension if it is an NLM. Because many NLM files require user input to unload, their unloading cannot be automated.

For Linux and Solaris, only enter the name of the process; you should not enter any path information. All processes running on the machine by that name will be stopped.

- ◆ **Wait for this process to finish running before continuing:** You can select this option for a process that unloads itself. By default, this option is deselected.

If an NLM requires intervention to unload, you must remember to unload it manually before trying to install the Distribution.

Start/Stop Windows Service Action

This action works for Windows services only, as illustrated in [Figure 3-17](#) and [Figure 3-18](#):

Figure 3-17 Properties of Distribution to See MSI Type Dialog Box

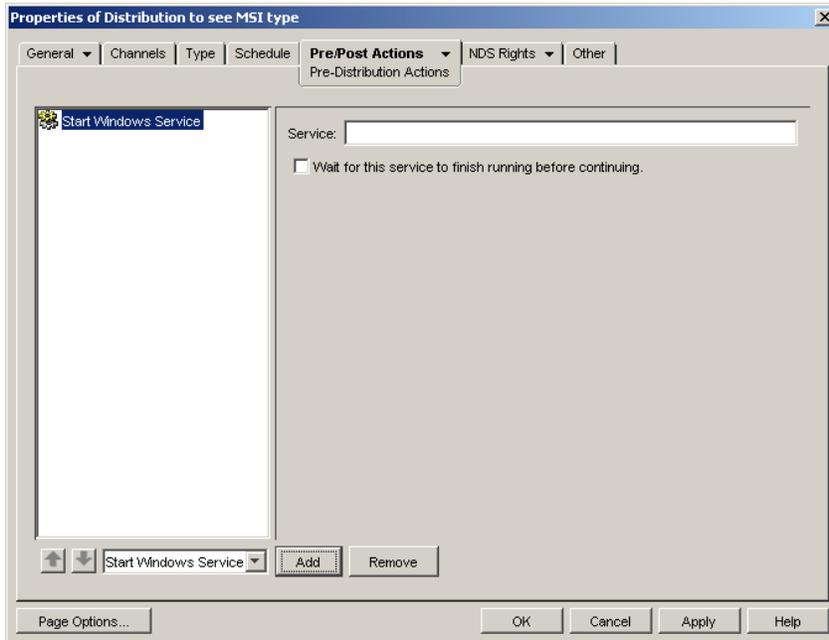
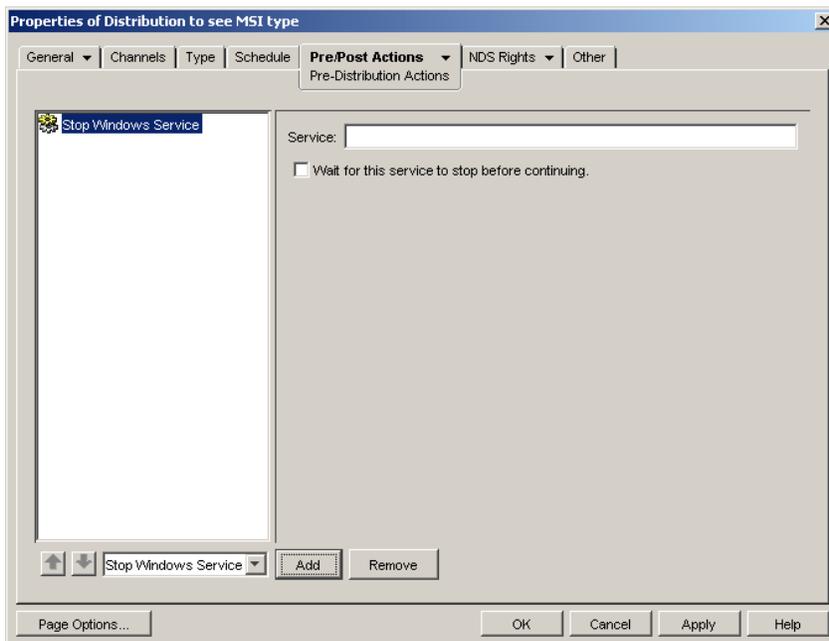


Figure 3-18 Properties of Distribution to See MSI Type Dialog Box



To add a Start/Stop Service action, select the option in the drop-down box and click the Add button. Then fill in the fields:

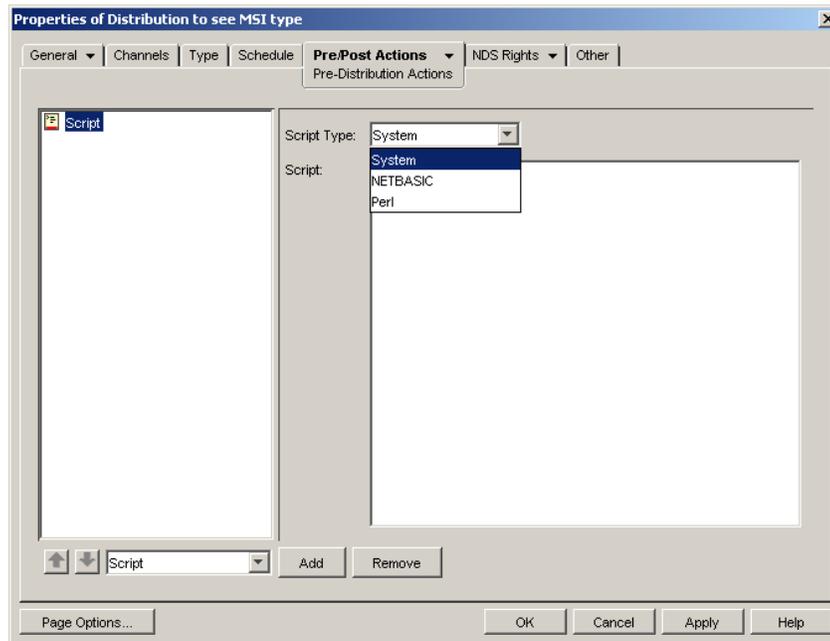
- ◆ **Service:** This must be the exact service name.
- ◆ **Wait for this service to finish running before continuing:** You can select this option for a service that starts or stops itself. By default, this option is deselected.

Script Action

This action works for Windows services, Java processes, and NLM processes.

You can run server scripts before installing the main Distribution files. Use the arrows to arrange the scripts' execution order.

Figure 3-19 Properties of Distribution to See MSI Type Dialog Box



To add a Script action, select the option in the drop-down box and click the Add button. The word “Script” defaults, which you must change to the script filename, including its full path. (Without the path, the script cannot be found to run it.)

Then fill in the fields:

- ◆ **Script type:** There are three script types: System, NetBasic, and PERL. The text you enter in the Script box must match the type you select.

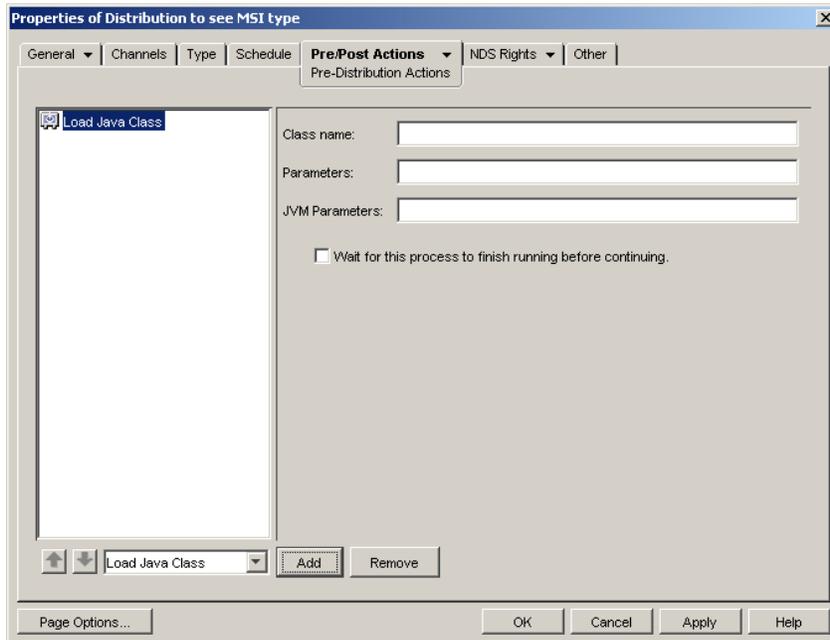
IMPORTANT: NetBasic is not supported on NetWare 6.5 servers.

- ◆ **Script:** Enter the text of the script.

WARNING: If a Distribution executes the script, processing done by the script cannot be undone by rollback.

Load Java Class Action

Figure 3-20 Properties of Distribution to See MSI Type Dialog Box



This action works for NetWare only.

To add a Load Java Class action, select the option in the drop-down box and click the Add button. Then fill in the fields:

- ♦ **Class name:** This must be the exact name. The `.class` extension is not necessary.

IMPORTANT: In order to load a Java class, `java.exe` or `jre.exe` must already be in the path on the server receiving the Distribution. Or, in this field, you can include the full path to the file.

- ♦ **Parameters:** Include any command line parameters for the Java application being run.
- ♦ **JVM parameters:** Include any parameters for the Java machine.
- ♦ **Wait for this process to finish running before continuing:** You can select this option for a Java application that terminates itself. It must terminate within 10 minutes, or the whole loading process fails. By default, this option is deselected.

3.4.7 Reassigning a Distribution to Another Distributor

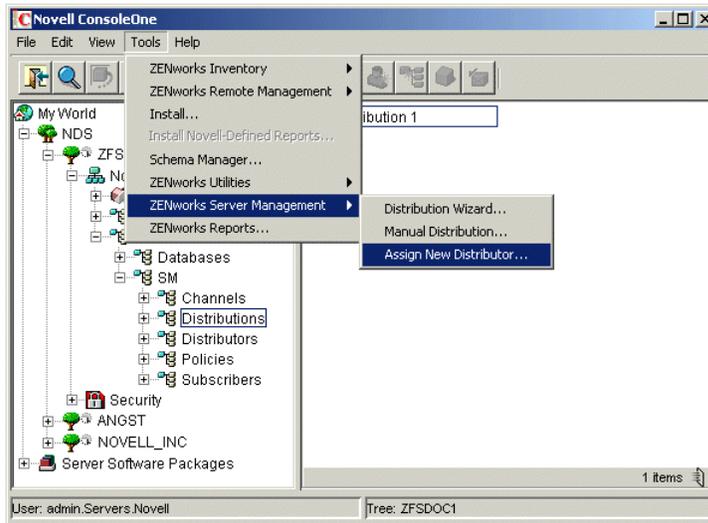
A single Distributor can service many Distributions, which could cause performance degradation on that Distributor's server. In version 7, there is a way to reassign a Distribution from one Distributor to another to balance the work load without needing to re-create the Distribution.

You can select one or more Distributions and reassign them to another Distributor.

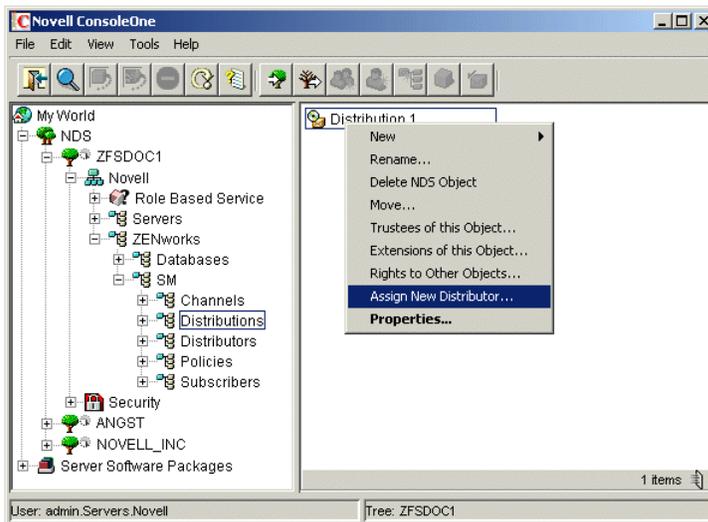
If you delete a Distributor object in ConsoleOne, you are asked if you want to reassign the Distributions that it services.

To reassign a Distribution to another Distributor:

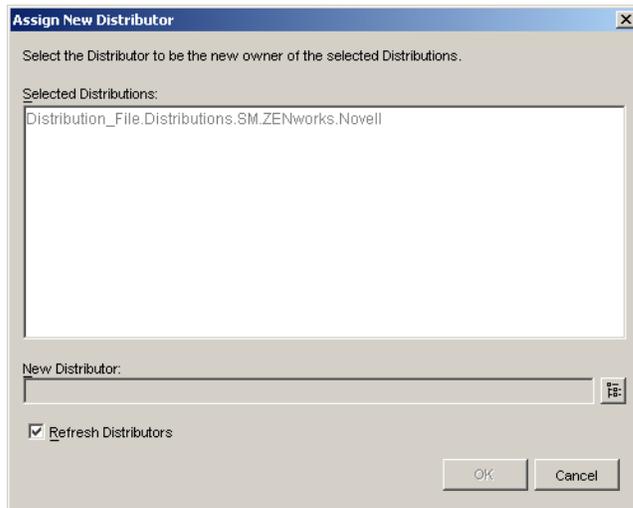
- 1 Determine which Distributions you want to reassign to another Distributor.
- 2 In ConsoleOne, do one of the following:
 - ♦ Select one or more Distribution objects, click *Tools*, click *ZENworks Server Management*, then click *Assign New Distributor*.



- ♦ Select one or more Distribution objects, right-click the selected objects, then click *Assign New Distributor*.



The following dialog box is opened when using either of the above options:



The Distributions you selected are listed in the *Selected Distributions* list.

If you want to change the list, you must click *Cancel* and reselect the Distribution objects.

- 3 In the *New Distributor* field, browse for the Distributor object that you want to be the new owner of these Distributions.

IMPORTANT: Any files on the current Distributor server's file system that belong in the Distribution must be copied or moved to the new Distributor server's file system, using the identical full path. This is covered in [Step 7](#).

- 4 If you want the Distributions to be built by the new Distributor owner as soon as you've finished reassigning them, select the *Refresh Distributors* check box.

The new Distributor is refreshed upon exiting this process (see [Step 5](#)), so that it immediately recognizes its new Distributions.

IMPORTANT: If you have files to copy, such as for the File, MSI, and Desktop Application types of Distributions, you should wait to refresh the new Distributor until after you have copied or moved the files for those Distributions to the new Distributor server's file system, this task is accomplished in [Step 7](#) through [Step 12](#).

- 5 Click *OK* to transition the Distribution objects to the new owner.
- 6 To make the old Distributor aware that it no longer has the Distributions that were reassigned, right-click the old Distributor's object, then click *Refresh Distributor*.

IMPORTANT: The reassignment tool in ConsoleOne only reassigns the eDirectory objects. Therefore, for File or MSI Distributions, the files contained in those Distributions reside on the old Distributor's file system. These files need to be moved to the new Distributor's file system so that the new Distributor has access to them for building these File or MSI types of Distributions. This is covered in [Step 7](#).

For Desktop Application Distributions, you need to review the Application objects to determine which files contained on the old Distributor's file system need to be moved to the new Distributor's file system. This is covered in [Step 10](#).

- 7 If a Distribution is a File or MSI type, do the following:

- 7a** In ConsoleOne, right-click the Distribution object for the Distribution that you want to reassign, then click *Properties*.
 - 7b** Select the *Type* tab.
 - 7c** In the *Files to be distributed* list, note all of the files or directories to be distributed, including their full paths.
 - 7d** Exit the Distribution object.
- 8** Using your file location notes and file management software (such as Windows Explorer), copy or move all of the Distribution's files from the current Distributor server's file system to the file system of the Distributor server that is the new owner of the Distribution.
- The full paths and filenames must exactly match between both of the Distributor servers' file systems. If you do not make the paths identical between the old and new Distributor servers, you need to edit the Distribution's properties to match the newer paths.
- 9** Repeat **Step 7** and **Step 8** for each Distribution to be reassigned.
 - 10** If a Distribution is a Desktop Application type, do the following:
 - 10a** In ConsoleOne, right-click the Distribution object for the Distribution that you want to reassign, then click *Properties*.
 - 10b** Select the *Type* tab.
 - 10c** Note which Application objects are in the Distribution, then note the `.fil` files for each Application object, including their full paths.
 - 10d** Exit the Distribution object.
 - 11** Using your file location notes and file management software (such as Windows Explorer), copy or move all of the Distribution's Application object files from the current Distributor server's file system to the file system of the Distributor server that is the new owner of the Distribution.
- The full paths and filenames must exactly match between both of the Distributor servers' file systems. If you do not make the paths identical between the old and new Distributor servers, you need to edit the Application object's properties to match the newer paths.

IMPORTANT: Although you normally have automatic temporary file clean-up for this Distribution, the temporary files for the Distribution being reassigned must be cleaned up manually from the old Distributor's server.

- 12** Repeat **Step 10** and **Step 11** for each Distribution to be reassigned.
- 13** If you did not elect to refresh the Distributors immediately, and you want the new Distributor to now recognize its new Distributions, right-click the new Distributor's object, click *Refresh Distributor*.

The previous Distributor no longer attempts to build the transitioned Distributions. The Distributor that now owns the Distributions is the one to build and send them, according to the Build and Send schedules.

3.4.8 Deleting a Distribution

If you delete a Distribution object, you must immediately refresh the Distributor that owned the Distribution; otherwise, the following can happen:

- ♦ When the Build schedule fires, the Distributor tries to build a Distribution that it thinks still exists, causing an error.

- ◆ In iManager, if you select the Distribution Information option for the deleted Distribution, the Distributor receives a 601 null-pointer error.

By immediately refreshing the Distributor, you prevent both of these errors from occurring, because:

- ◆ The Distributor reads eDirectory when it is refreshed and no longer knows of the deleted Distribution.
- ◆ The Distribution Information option for the deleted Distribution is no longer available in iManager.

3.4.9 Removing a Distribution Object - Auto Removal of Temporary Files

Previously, when you deleted a Distribution or Channel object, removed a Distribution or Subscriber from a Channel, or in some way caused one or more Distributions to no longer be associated with one or more Subscribers, the Distributions' temporary files remained on the Subscriber servers, and you had to find them and delete them manually to recover disk space.

In version 7, when a Distributor refreshes, the temporary files of the Distributions that have been removed (either deleted or removed from a Channel) are automatically deleted from Subscribers to free up disk space.

What Causes Temporary Distribution Files To Be Cleaned Up

A Distribution's temporary files are removed from a Subscriber server's file system when:

- ◆ The Distribution object is deleted
- ◆ The Channel object hosting the Distribution is deleted
- ◆ The Distribution is removed from the Channel
- ◆ The Subscriber is unsubscribed from the Channel

What the Distributor Does

When a Distributor refreshes, it determines whether any servers (including parent Subscribers in its routing hierarchy) still need to receive any of its Distributions. Where it is found that a Distribution is no longer needed, the Distributor notifies the Subscribers (including parent Subscribers) to clean up that Distribution's temporary files.

If a Distributor cannot contact a Subscriber or has not received a successful deletion reply, it sends another notification to that Subscriber the next time the Distributor refreshes. Therefore, the Refresh schedule determines how often a Subscriber is notified to clean up a deleted Distribution's temporary files.

A Distributor tries five times to notify a Subscriber to clean up a Distribution. If unsuccessful, the Distributor ceases notifying the Subscriber. Then, the temporary files on the Subscriber server can only be cleaned up manually.

What the Subscriber Does

When a Subscriber receives a notification to remove a Distribution's temporary files, it first determines whether the Distribution to be cleaned up is in the process of being received, sent, or extracted by the Subscriber server. If it is not, the Subscriber removes any forwarding or extraction

events that are pending and deletes the Distribution's directory containing the temporary files. Then, the Subscriber notifies the Distributor of the removal so that the Distributor can keep track of which Subscribers have successfully complied.

Parent Subscribers are treated the same as end-node Subscribers for cleaning up Distribution files.

Clean Up of Temporary Distribution Files on the Distributor Server

When the Distributor determines that a Distribution object is deleted from eDirectory, the Distribution's version directories (not the Distribution's directory) are automatically deleted from the Distributor's working directory.

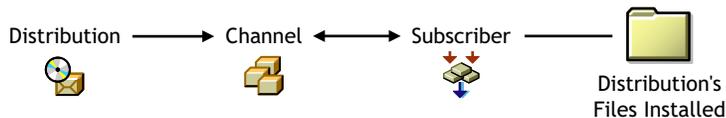
After the Distributor has determined that all notified Subscribers have successfully deleted the Distribution directories from their file systems, the Distributor then deletes the Distribution's directory from its file system.

When Subscribers Must Wait to Clean Up Temporary Distribution Files

Temporary Distribution files cannot be deleted from a Subscriber's file system until the association between the Distribution and the Subscriber is broken. For example:

- ◆ When a Distribution is listed in the Channel where the Subscriber is subscribed, the Distribution's files can be received and extracted on the Subscriber server:

Figure 3-21 Temporary Distribution File Cleanup: A



- ◆ If the Subscriber is no longer subscribed to the Channel, or the Distribution is no longer listed in the Channel, the Distribution's temporary files can be deleted from the Subscriber server:

Figure 3-22 Temporary Distribution File Cleanup: B

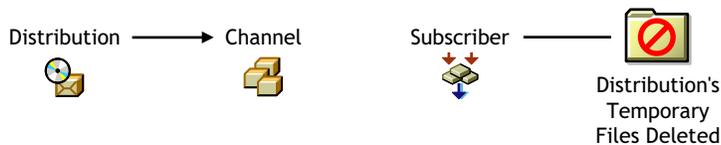
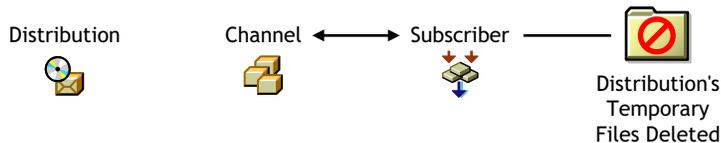


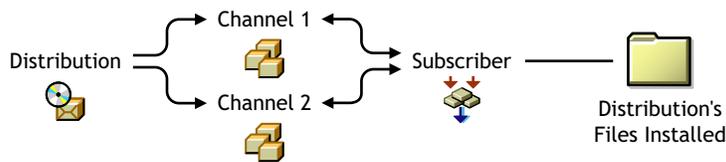
Figure 3-23 Temporary Distribution File Cleanup: C



However, if a Distribution and a Subscriber are associated through multiple Channels, the Distribution's temporary files are not deleted from the Subscriber's file system until both the Distribution and Subscriber objects are no longer associated through any Channel. For example:

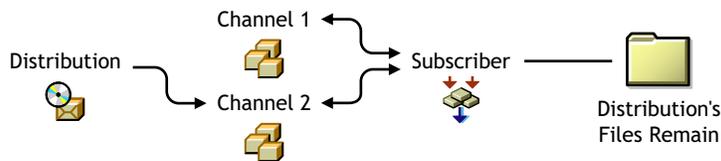
- ◆ When a Distribution is listed in two different Channels and the Subscriber is subscribed to both Channels, the Distribution's files can be received and extracted on the Subscriber server:

Figure 3-24 Temporary Distribution File Cleanup: D



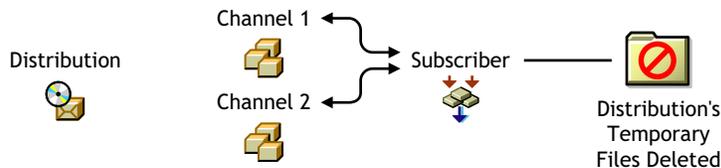
- ◆ When the Distribution is removed from one of the Channels, the Distribution's files can still be received and extracted on the Subscriber server:

Figure 3-25 Temporary Distribution File Cleanup: E



- ◆ When the Distribution is removed from both Channels, the Distribution's temporary files can be deleted from the Subscriber server:

Figure 3-26 Temporary Distribution File Cleanup: F



When a Parent Subscriber Can Remove the Temporary Distribution Files

There are two possibilities for when a parent Subscriber can have Distribution files on its server:

- ◆ When both the parent Subscriber and the end-node Subscriber are subscribed to the same Channel to receive its Distributions. The parent Subscriber passes on the Distributions and also extracts them for itself.
- ◆ When the parent Subscriber is not subscribed to the Channel that the end-node Subscriber is. The parent Subscriber only passes on the Channel's Distributions.

In both cases, the following rules apply to when Distribution files can be cleaned up from a parent Subscriber's server:

- ◆ If the parent Subscriber unsubscribes to the Channel, the Distributions' files are not deleted from the parent Subscriber's server, so that it can continue to forward those Distributions to the end-node Subscriber server.
- ◆ Only after the end-node Subscriber unsubscribes from the Channel is the parent Subscriber able to clean up the Distributions' files from its server.

3.4.10 Handling Orphaned Distributions

The following sections explain how to handle the Distributions of a deleted Distributor object:

- ◆ [“Orphaned Distributions” on page 140](#)

- ◆ [“Cleaning Up Orphaned Distributions” on page 140](#)
- ◆ [“Re-creating Deleted Distributions” on page 140](#)

Orphaned Distributions

Because Distributions belong exclusively to their Distributors, you cannot build and send those Distributions if you delete a Distributor object from eDirectory. The Distributions associated with the deleted Distributor become orphaned and are no longer usable.

Any orphaned Distributions that have already been sent and extracted before you delete the Distributor object are usable by the Subscriber servers where they were extracted. However, these servers no longer receive updated versions of the orphaned Distributions.

You can still see the orphaned Distribution objects in eDirectory, but no current or future Distributor object can be associated with these orphaned Distribution objects.

Cleaning Up Orphaned Distributions

For all Distribution types, you can delete the Distribution’s directories on the Subscriber servers’ file systems for all orphaned Distributions. We recommend that you delete the Distribution’s directories for any Distributions that you intend to re-create.

For most Distribution types, deleting the orphaned Distributions’ directories is all you need to do in order to clean up for management and disk space conservation purposes. These Distribution types are:

- Desktop Application
- File
- FTP
- HTTP
- RPM

However, for the Policy Package and Software Package Distribution types, you might need to undo the processes that the Distributions initiated when they were extracted and installed.

For example, a Policy Package Distribution might require that you use iManager to remove the policies that the Distribution set for the server. For more information, see [Step 5](#) under [Section 2.4.3](#), [“Managing the Policy/Package Agent,” on page 80](#).

Re-creating Deleted Distributions

You need to re-create each orphaned Distribution that you want to continue to use. You can do this using an existing Distributor object, or after you install a new Distributor.

After you have re-created a Distribution, all Channels previously associated with the orphaned Distribution need to be associated with the newly created Distribution.

In re-creating the Distributions, you can use the configuration information from the orphaned Distribution objects. When you no longer need the orphaned Distribution objects, you can delete them and they no longer display on the Distributions tab of the Channel object.

3.4.11 Manually Importing and Exporting Distributions

Exporting and importing are useful for:

- ◆ Sending a large Distribution to Subscriber servers that are across a slow WAN link from the Distributor server.
- ◆ Sending a large Distribution to a parent Subscriber server that is across a slow WAN link, then having that parent pass the Distribution on to its subordinate Subscribers on its side of the WAN.
- ◆ Archiving Distributions, and later importing them when and where they are needed again.

The following sections provide information on exporting and importing Distributions:

- ◆ [“Understanding the Exporting and Importing Processes” on page 141](#)
- ◆ [“Setting Up Specialized Schedules” on page 141](#)
- ◆ [“Exporting a Distribution” on page 142](#)
- ◆ [“Importing a Distribution” on page 142](#)

Understanding the Exporting and Importing Processes

You can manually export a Distribution from a Distributor server by writing to a media source, such as a floppy disk, ZIP disk, CD, or DVD, then you can import it from that media to a Subscriber server.

The export process copies Distribution information to a UNC path or drive mapping, such as a hard drive, floppy disk, or ZIP disk. From the copy on the hard drive, you can then burn the information onto a CD or DVD.

The Distribution information includes the Channel and Distribution data from their eDirectory objects, and the content of the Distribution’s file (including all deltas). The Distribution information is copied to a `filename.ted` file that you name when running the Manual Distribution Wizard. You should use the `.ted` extension with the filename. You should also use a descriptive filename so that you can recognize the Distribution when reviewing the media content.

When the exported `.ted` file is imported, the eDirectory object information and the Distribution’s content are used to create the Distribution on the Subscriber server’s file system. Thereafter, deltas of the Distribution can be sent over the wire, because they are usually much smaller than the original Distribution that was exported and imported.

Distributions can only be exported and imported within the same tree where the associated Channels are known to all Distributors and Subscribers involved.

Setting Up Specialized Schedules

Depending on when you want imported Distributions to be extracted, you might want a different set of schedules set up before exporting the Distribution.

For example, if you want the exported Distribution to be extracted at different times by different Subscribers where it is imported, then:

- 1 Set the build schedule for the Distribution to be exported to *Immediate*.
- 2 Add the Distribution to a Channel with a Send schedule set to *Never*.

This prevents Subscribers that have not yet had the Distribution manually imported to them from receiving a Channel's notice to trigger extraction of the yet-to-be-received Distribution.

- 3 Add all of the Subscribers where the Distribution is to be imported to the Channel you used in [Step 2](#).
- 4 Refresh the Distributor that owns the Distribution to be exported.
- 5 After the Distribution has been built, continue with [“Exporting a Distribution” on page 142](#).

If you do not need a specialized schedule, you can just follow the instructions in the next two sections, which assume that existing schedules are acceptable.

Exporting a Distribution

- 1 In ConsoleOne, click *Tools*, then click *Manual Distribution* to start the Manual Distribution Wizard.
- 2 Click *Export*, then click *Next*.
- 3 Select a Channel, select one Distribution from that Channel, then click *Next*.
This Channel's ID is retained in the `.ted` file for use when importing the Distribution.
- 4 For the Distribution, provide a path (UNC or drive mapping) and filename (descriptive for identifying which Distribution is on the media), then click *Next*.
The filename should have `.ted` as its extension.
- 5 If you are satisfied with the summary, click *Finish*.
The full Distribution is saved as a `.ted` file to the path that you specified.
- 6 If your path was to a hard drive, you can now burn the `.ted` file to a CD or DVD.

Importing a Distribution

- 1 In ConsoleOne, click *Tools*, then click *Manual Distribution* to start the Manual Distribution Wizard.
- 2 Click *Import*, then click *Next*.
- 3 Provide the path and filename to the `.ted` file, then click *Next*.
This is the `.ted` file that you exported to a media source.
- 4 Select parent Subscribers in the top box and individual Subscribers in the bottom box, then click *Next*.

If you select a parent Subscriber that is in the routing hierarchy, all of the Subscribers below it in the hierarchy have the imported Distribution passed on to them, but only if they are already subscribed to the Distribution's Channel.

The Subscribers displayed in the bottom box are those who are currently subscribed to the Distribution's Channel. The heading displays the Channel that is associated with the Distribution being imported. This information is contained in the `.ted` file being imported.

External Subscribers are not listed in the bottom box because they cannot receive manual Distributions.

- 5 If you are satisfied with the summary, click *Finish*.
The Distribution is copied from the media source you specified and placed in the working directories of the selected Subscribers. The Channel and Distribution objects' information is written to eDirectory.

At this point, imported Distributions are viewable in Remote Web Console in iManager, but not in Tiered Distribution View or Subscriber Distribution View. The next two steps take care of this.

- 6** If you set up specialized schedules for the imported Distribution (see “[Setting Up Specialized Schedules](#)” on page 141), restart the Server Management process on each Subscriber server where it was imported; otherwise, skip to [Step 7](#).

The Distribution is extracted on the Subscriber servers according to their individual Extract schedules. After extraction, you can view the Distribution’s information in iManager.

- 7** To make Distributors recognize that their Subscribers have manually received a new Distribution:
 - 7a** Under the ZENworks Server Management role in iManager, click *Remote Web Console*.
 - 7b** Identify the Distributor owning the imported Distribution in either of the following fields:
 - IP Address or DNS Name
 - Distributor, Subscriber, or Server Object Name
 - 7c** Click *OK*.
 - 7d** In the *Display* field, select *Tiered Electronic Distribution*.
 - 7e** Click the *Channels* tab, then select *Distribute Channel*.
 - 7f** Click the Channel associated with the imported Distribution, then click *OK*.

The Distributor begins to send the Distribution listed in the Channel to the Subscribers, but the Subscribers reply that they already have the Distribution, then begin to extract the imported Distribution.

If a Subscriber is a parent Subscriber that needs to pass the imported Distribution on to subordinate Subscribers, it does so when the Distribution’s Channel starts.

3.4.12 Using the Distribution Wizard

Server Management provides the Distribution Wizard to help you learn the process involved in creating and sending a Distribution. You can use this wizard to create and send either a File or FTP Distribution.

To use the Distribution Wizard:

- 1** In ConsoleOne, select the container where you want the Distribution object created, click *Tools*, then select *Distribution Wizard*.
- 2** Review the information on the Introduction page, then click *Next*.
- 3** On the Distributor Selection page, browse for and select the Distributor that owns this File or FTP Distribution, then click *Next*.
- 4** On the Subscriber Selection page, click *Add*, browse for the Subscribers to receive this Distribution, click *Select*, click *OK*, then click *Next*.
- 5** On the File Source page, select the file source (the Distributor’s file system, or a remote FTP site), then click *Next*.
- 6** On the Destination Volume or Drive page, select an option and fill in its field, then click *Next*.

Use the same volume or drive for all Subscribers: If each target Subscriber is to have the exact same volume or drive available, select this option and provide the volume label or drive letter.

Use a variable for the volume or drive: If your target Subscribers are using different paths (for example you have NetWare, Windows, Linux, and Solaris Subscriber servers), you can provide a variable value. This value must be defined on each Subscriber in order to receive the Distribution.

- 7 On the Additional Destination Directories page, provide any additional path information for the target Subscriber servers, then click *Next*.

Your path information is displayed under the *Data Will Be Placed In Path* heading as you enter it. Use this information to verify that the path is valid before continuing.

- 8 On the File Selection From Distributor Server page, click *Add*, browse for the files or directories to be included, click *Select*, click *OK*, then click *Next*.

You are browsing the Distributor's file system, not the local machine's.

Repeat clicking *Add* until you have all of the files and directories you want in this Distribution.

- 9 On the Distribution Name and Context page, fill in the fields, then click *Next*.

Distribution name: Provide a unique name for the Distribution.

Context: Browse for and select the container where you want the Distribution object to be created.

- 10 On the Additional Options page, select or deselect the options as applicable, then click *Next*.

The following options are all selected by default:

Copy the Distributor's security certificate to all Subscribers: This is necessary for the Subscriber to be able to receive and extract this Distribution. This might not be necessary if you run the wizard again with the same Distributor and Subscribers.

Verify that all Subscribers are up and running: If you want to make sure your target Subscribers can receive this Distribution, select this option.

Notify the Distributor to read eDirectory for new information: This causes the Distribution to be built immediately.

- 11 On the Summary page, review the steps that are take by the Distribution Wizard, then click *Finish* to create the Distribution.

Information is displayed as the Distribution is created and sent.

- 12 To review the log file, select *Yes* when prompted.

If you select *Yes*, you can review the log file. Click *Close* to exit the log window and the Distribution Wizard.

If you select *No*, the Distribution Wizard is exited.

3.5 Channels

The following sections provide concepts and instructions for the Channel object:

- ♦ [Section 3.5.1, "Understanding Channels," on page 145](#)
- ♦ [Section 3.5.2, "Creating and Configuring Channels," on page 146](#)
- ♦ [Section 3.5.3, "Forcing a Channel To Be Sent," on page 147](#)

3.5.1 Understanding Channels

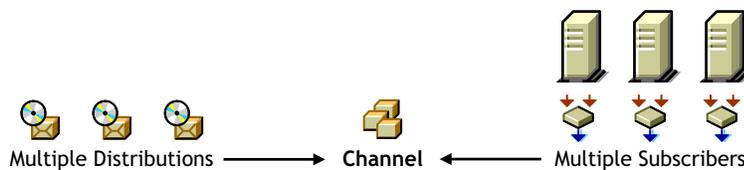
The Channel object (TED Channel) contains a list of Distributions associated with it and Subscribers subscribed to it.

- ♦ [“Functional Relationship with Other Tiered Electronic Distribution Objects” on page 145](#)
- ♦ [“Channel Description” on page 145](#)
- ♦ [“Scheduling” on page 145](#)
- ♦ [“Subscriptions to Channels” on page 145](#)

Functional Relationship with Other Tiered Electronic Distribution Objects

[Figure 3-27](#) illustrates a Channel’s relationship with Distributions and Subscribers:

Figure 3-27 Channel Relationship with Distributions and Subscribers



The Distributions are listed in the Channel, and the Subscribers subscribe to the Channel.

Channel Description

Distributors can list one Distribution in multiple Channels, and multiple Distributors can list their Distributions in the same Channel.

You can have as many Channels as you want. Channels do not hold the actual Distributions, only a reference to them. There is no limit to the number of Distribution references a Channel can send. The practical limit is how many Distributions you want to track per Channel.

Scheduling

A Channel’s Send schedule determines when a Distribution are sent from the Distributor to its Subscribers.

A Channel can be active or inactive to control when its Distributions are sent.

For information on how time zones can affect scheduling between a Channel and its associated Distributors and Subscribers, see [“Scheduling Tiered Electronic Distribution Objects in Different Time Zones” on page 334](#).

Subscriptions to Channels

Channels can be subscribed to by multiple Subscribers.

To receive a Distribution, a Subscriber must subscribe to the Channel where that Distribution is listed. However, a Subscriber receives all of the Distributions listed in that Channel, which means they are applied to the Subscriber server when they are extracted.

3.5.2 Creating and Configuring Channels

The following sections provide you with the steps to create and configure the Tiered Electronic Distribution objects with ConsoleOne.

Do the following in order for each Distributor:

- ◆ [“Determining the Channel Names” on page 146](#)
- ◆ [“Creating the Channel Objects” on page 146](#)
- ◆ [“Configuring the Channels” on page 146](#)

Determining the Channel Names

In naming Channels, use a descriptive method. For example:

```
VirusProtect  
VProtectPatterns  
VirusProtection  
NW51patch4  
NW6patch1  
AUTOEXECNCF000326
```

You can manage your Channels more easily by:

- ◆ Using names that are purpose oriented
- ◆ Using a similar name for the Channel and its Distributions

Continue with [“Creating the Channel Objects” on page 146](#).

Creating the Channel Objects

Channels are used to group Distributions and establish a schedule for passing a Distributor’s Distributions to Subscribers that are subscribed to the Channel. A Channel can have Distributions from many Distributors. A Channel can be subscribed to by many Subscribers.

To create a Channel object:

- 1 In ConsoleOne, select a container object to hold the Channel object, click *File > New > Object*, then select *TED Channel*.
- 2 Provide a name for the Channel object and click *OK*.
- 3 Create as many Channel objects as needed to group Distributions by type and/or send schedule.
- 4 Continue with [“Configuring the Channels” on page 146](#).

Configuring the Channels

You need to configure a Channel object before you can begin using it.

Not all properties associated with the Channel object are required. Required objects are noted; all others are optional.

To configure the Channel object:

- 1 In ConsoleOne, right-click the Channel object, then click *Properties*.
- 2 Select the *General* tab and fill in the fields:

Active: Select the check box to enable the Channel to pass on its Distributions.

Description: Provide a useful description, such as what Distributions the Channel is associated with.

3 Select the Distributions tab, then click *Add* to add Distributions.

Distributions: A list of Distributions that are associated with this Channel. For information on creating Distribution packages, see [Section 3.4, “Distributions,” on page 110](#).

4 Select the *Subscribers* tab, then click *Add* to add Subscribers to the Channel.

Subscribers subscribed to this Channel: A list of Subscribers and External Subscribers that are subscribed to this Channel.

5 Select the *Schedule* tab, then select a schedule for when to distribute the Channel’s Distributions.

For information on available schedules, see [Chapter 8, “Scheduling,” on page 317](#).

3.5.3 Forcing a Channel To Be Sent

If you want to send all of the Distributions in a Channel outside of Channel’s the normal Send schedule, you can manually force the distribution process.

Assuming that a new Distribution has been built and the Channel’s Send schedule is not ready to fire, do one of the following to force a Channel to be sent:

- ◆ Using the ZENworks Server Management role in iManager, click *Edit TED Object*, browse for and select the Channel, click *OK*, then click *Distribute Channel*.
- ◆ In ConsoleOne, you have a two-step process:
 - a. Select the Channel object, click *Properties*, select the *Schedule* tab, select *Run Immediately*, click *OK*, right-click the Distributor object, then click *Refresh Distributor*.
 - b. After the Distribution has been sent, to reverse the changes made in Step a, select the Channel object, click *Properties*, select the *Schedule* tab, select the schedule that the Channel previously had, then click *OK*.

As soon as a Subscriber receives an entire Distribution, it extracts according to the Subscriber’s Extract schedule.

3.6 Subscribers

The following sections provide concepts and instructions for the Subscriber object:

- ◆ [Section 3.6.1, “Understanding Subscribers,” on page 148](#)
- ◆ [Section 3.6.2, “Creating Subscribers,” on page 149](#)
- ◆ [Section 3.6.3, “Configuring Subscribers,” on page 150](#)
- ◆ [Section 3.6.4, “Updating Subscriber Configurations,” on page 153](#)
- ◆ [Section 3.6.5, “Associating Subscribers with Channels,” on page 154](#)
- ◆ [Section 3.6.6, “Deleting Subscriber Objects That Are Part of a Distributor’s Routing Hierarchy,” on page 155](#)

3.6.1 Understanding Subscribers

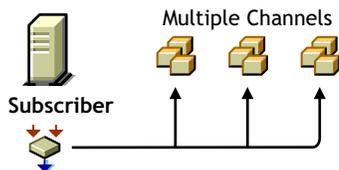
The Subscriber object (TED Subscriber) is an eDirectory object that defines the properties for the Subscriber.

- ◆ “Functional Relationship with Other Tiered Electronic Distribution Objects” on page 148
- ◆ “Subscriber Description” on page 148
- ◆ “Scheduling” on page 149
- ◆ “Subscribing to Channels” on page 149
- ◆ “Parent Subscribers” on page 149
- ◆ “Special Character Handling” on page 149

Functional Relationship with Other Tiered Electronic Distribution Objects

Figure 3-28 illustrates a Subscriber’s relationship with the Channels:

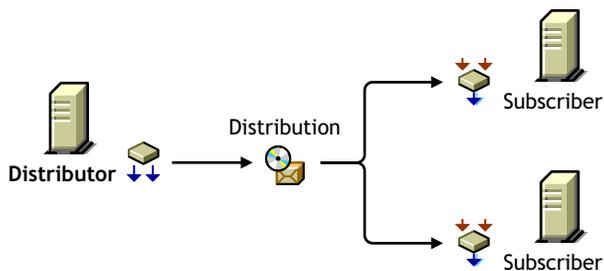
Figure 3-28 A Subscriber that Subscribes to Multiple Channels



The Subscriber subscribes to the Channels.

Figure 3-29 illustrates the Subscriber’s relationship with Distributors and Distributions:

Figure 3-29 Multiple Subscribers Can Receive the Same Distribution from a Distributor.



Subscriber Description

The Subscriber is a service that receives and extracts Distributions to obtain the software, files, or policies it needs.

Any server where you want to distribute applications, files, or policy packages must have the Subscriber software installed and a Subscriber object in the eDirectory tree. The Subscriber object can be in a different tree than the server’s NCP™ server object, because IP addresses or DNS names are used for moving Distribution files to the Subscriber servers.

Distributions are copied to the Subscriber server’s hard drive. The Subscriber receives the Distributions and extracts them to install the software, files, or policies.

Scheduling

A Subscriber's Extract schedule determines when it can extract its Distributions.

For information on scheduling, see [Chapter 8, "Scheduling," on page 317](#).

Subscribing to Channels

Subscribers can subscribe to a Channel to receive all of the Distributions listed in that Channel. A Subscriber object's properties lists the Channels it is subscribed to.

Subscribers can receive Distributions from multiple Distributors because:

- ◆ Multiple Distributors can list their Distributions in the same Channel
- ◆ Subscribers can subscribe to multiple Channels

Parent Subscribers

Subscribers can be parent Subscribers, which are proxies for the Distributor to pass Distributions to other Subscribers. This helps the Distributor by providing load-balancing for sending Distributions to many Subscribers.

The Subscriber object's properties lists the parent Subscriber through which it receives all of its Distributions. A Subscriber can receive its Distributions directly from the Distributor if it does not have a parent Subscriber and is not listed in the Distributor's routing hierarchy.

Parent Subscribers can also be used to bridge WAN links to ensure that Distribution packages are sent across WAN links a minimum number of times.

Special Character Handling

Syntax differences (such as characters that are invalid to a platform) are now handled for each supported platform. For invalid characters, the agent properly gathers all files, regardless of platform of the Distributor server. The Subscriber server detects whether files in the Distribution package include invalid characters and ignores or skips files during extraction. Skipped files are logged. Previously, the whole Distribution would fail to extract and be installed.

Linux and Solaris support characters in file and directory names that NetWare and Windows do not recognize.

3.6.2 Creating Subscribers

Subscribers must be created by installing their software and eDirectory objects using the *ZENworks 7 Server Management with Support Pack 1 Program CD*. For more information, see "[Installation on NetWare and Windows Servers](#)" in the *Novell ZENworks 7 Server Management Installation Guide*.

If a Subscriber object is inadvertently deleted, you can re-create it in ConsoleOne. However, the Revision Number of the new Subscriber object will be less than its Revision Number in the `ted.cfg` file. Therefore, the Subscriber cannot accept any updates to its configuration, because the lower Revision Number causes it to assume that the configuration data is older than what it has. To resolve this problem, delete the `ted.cfg` file on the Subscriber server, and the next time a Distribution is sent to the Subscriber, a new configuration is accepted, and a new `ted.cfg` file created.

3.6.3 Configuring Subscribers

Subscriber objects are automatically created when you install the Subscriber software to a server.

Not all properties associated with the Subscriber object are required. Required objects are noted; all others are optional.

To configure the Subscriber object's properties:

- 1 In ConsoleOne, right-click the Subscriber object, then click *Properties*.
- 2 Click *General > Settings* and fill in the following fields:

Use policy: Select to use the effective policy if you want to use the values set in the Tiered Electronic Distribution policy. This field is displayed if a Tiered Electronic Distribution policy has been created, distributed to the Subscriber server, extracted by the Policy/Package Agent, and enforced on the server.

If you select this option, the rest of the fields are dimmed and the policy settings are used instead. The current policy is displayed in parentheses.

Input rate: The rate Distributions are received. The default is the maximum that the connection can handle. This rate is used to control a Subscriber server's use of narrow bandwidth links.

Output rates based upon Distribution's priority: Sets the default output rate to minimize network traffic for Tiered Electronic Distribution objects. This determines the send rate for parent Subscribers to its subordinate Subscribers. The default value is the maximum that the connection can handle. Blank means that bandwidth is taken from third-party applications.

There are three output priorities where you can specify a rate:

- ♦ **High priority:** These Distributions are sent before any Medium or Low priority Distributions.
- ♦ **Medium priority:** These Distributions are sent after all High priority and before any Low priority Distributions.
- ♦ **Low priority:** These Distributions are sent after all High and Medium priority Distributions.

For more information, see [Section 3.4.5, "Prioritizing Distributions," on page 126](#).

Maximum concurrent Distributions to send: Specifies the maximum number of distribution threads that can be running concurrently for sending on Distributions. The default value is unlimited (a blank field).

This applies only to parent Subscribers that pass on Distributions to subordinate Subscribers.

Connection time-out: Specifies the number of seconds a Subscriber waits for a response from a Distributor (receiving) or a Subscriber (sending) before ending the connection. If a connection is ended during sending or receiving, the send does not start again until the next time the Channel schedule starts. It then picks up where it left off.

The default value is 300 seconds (five minutes). The available range in seconds is 1 to 60,000. You should make this setting a reasonable time to wait for a response from one node to another.

This interval should be increased on slow or busy links where longer delays are frequent.

Working directory: Specifies the directory to be used by the Distribution. It contains Distributions, persistent status, and temporary working files. The default path is:

- ♦ **NetWare:** `sys:\zenworks\pds\ted\sub`

IMPORTANT: The default volume is sys: on NetWare servers. We recommend that you do not use the sys: volume because the content of this directory can become quite large.

- ♦ **Windows:** `c:\zenworks\pds\ted\sub`
- ♦ **Linux and Solaris:** `/var/opt/zenworks/zfs/pds/ted/sub`

For more information on the working directory, see [Section 3.12, “Working Directories,” on page 186](#).

Parent Subscriber (optional): Specifies a parent Subscriber from which Distributions are received.

This field is where you can enable efficient distribution from a Distributor to its Subscribers. The routing information in a Distributor object’s properties accounts only for parent Subscribers (the tiered distribution model). End-node Subscribers (most of the Subscribers in your tree) should not be listed there.

This field allows you to specify for each end-node Subscriber that it receives its Distributions via a specific parent Subscriber, instead of directly from the Distributor. This reduces the workload on the Distributor server, and provides the tiered distribution model for efficient sending of Distributions.

This field is also useful for allowing a parent Subscriber to send a Distribution to an External Subscriber’s server in another tree.

Disk space desired to be left free: Use this value to ensure there is enough free disk space for receiving Distributions. A Subscriber does not attempt to receive a Distribution if the disk space value set here is insufficient.

- 3 Click *General > Messaging* and fill in the following fields:

IMPORTANT: If this Subscriber is on the same server as a Distributor, entries in these fields are ignored. Only the Distributor’s messaging settings are used.

Use policy: Select to use the effective policy if you want to use the values set in the Tiered Electronic Distribution policy. This field is displayed if a Tiered Electronic Distribution policy has been created, distributed to the Subscriber server, extracted by the Policy/Package Agent, and enforced on the server.

If you select this option, the rest of the fields are dimmed and the policy settings for messaging are used instead. The current policy is displayed in parentheses.

Server console: Specifies the level of output messages to send to the Subscriber console on the server console.

For more information on the message notification levels, see [Section 3.11.5, “Minimizing Messaging Traffic,” on page 183](#).

SNMP trap: Specifies the level of messages to send via SNMP.

Log File: Specifies the level of messages to send to the log file.

Path and filename: You can specify a custom log file’s name and location for this Subscriber object. The default is:

- ♦ **NetWare:** `sys:\zenworks\pds\ted\dist\ted.log`

IMPORTANT: The default volume is sys: on NetWare servers. We recommend that you do not use the sys: volume because the log file can become quite large.

- ♦ **Windows:** `c:\zenworks\pds\ted\dist\ted.log`

- ♦ **Linux and Solaris:** `/var/opt/zenworks/zfs/pds/ted/dist/ted.log`

This is the same log file that the Distributor uses.

Delete log entries older than __ days: Log file entries for a Subscriber are deleted after they are older than the number of days specified. The default is six days.

E-mail: Specifies which level of messages to send via e-mail.

Users: Specifies e-mail users for notification.

Address attribute: Specifies e-mail addresses for notification.

You can add users or groups stored in eDirectory or provide the e-mail addresses for users who are not contained in eDirectory. The e-mail Address Attribute associated with an eDirectory user is the default attribute.

IMPORTANT: If you select e-mail as a method for receiving notification, be aware that additional network traffic can be created.

- 4 Click *General* > *Working Context* and browse for a working context.

This is the eDirectory context where the Subscriber creates the objects related to the Desktop Application Distributions it receives.

- 5 Select the *Schedules* tab, select a schedule, then fill in the fields:

Use policy: Select to use the effective policy if you want to use the values set in the Tiered Electronic Distribution policy. This field is displayed if a Tiered Electronic Distribution policy has been created, distributed to the Subscriber server, extracted by the Policy/Package Agent, and enforced on the server. If you select this option, the rest of the fields are dimmed and the policy settings for scheduling are used instead.

Schedule type: This schedule determines when the Subscriber extracts the Distributions.

For information on available schedules, see [Chapter 8, “Scheduling,” on page 317](#).

- 6 Select the *Channels* tab and fill in the fields:

- ♦ **Channels this Subscriber is subscribed to**

Lists the Channels the Subscriber is subscribed to.

Active: To activate a Channel for this Subscriber server so it can receive the Channel’s Distributions, click a Channel, then select the check box to enable it. To deactivate a Channel so that the Subscriber does not receive the Channel’s Distributions, deselect the check box to disable it.

Channel: Click Add to create a Channel. Click Details to edit a Channel.

- ♦ **Channels subscribed to through Subscriber Group memberships**

Lists the Subscriber Groups that the Subscriber is a member of, paired with which Channels the Subscriber is subscribed to by virtue of membership in a Subscriber Group.

These columns are for display only. The Details, Add, and Delete buttons do not apply.

Active: Indicates whether the Channel subscribed to is active.

Channel: Displays the Channel subscribed to through membership in a group.

Subscriber Groups: Displays the groups the Subscriber is a member of. You can sort the listing by clicking the column heading.

- 7 Select the *Variables* tab and fill in the fields:

Include policy: Select to use the effective policy if you want to use the values set in the Tiered Electronic Distribution policy. This field is displayed if a Tiered Electronic Distribution policy has been created, distributed to the Subscriber server, extracted by the Policy/Package Agent, and enforced on the server.

If you select this option, the variables specified in the Tiered Electronic Distribution policy are added to the list of variables. However, if there are duplicate variables, the variables in the Subscriber prevail.

Variable: Name of the variable. It should indicate how the variable is used. For example, WORKINGVOL.

Value: The value that the Subscriber uses when this variable is specified. For example, data:.

To ensure that extraction takes place, provide an absolute path to the Subscriber. For example, if the path is only the data: volume, make sure the colon (:) is included, because it is a necessary part of the full path.

Description: Describes how the variable is used. For example:

`Volume for the working directory.`

For information on variables, see [Section 9.6, “Using Variables to Control File Extraction,” on page 349.](#)

- 8 To include this Subscriber in a group, click *Group Membership*, click *Add*, browse for a Subscriber Group object, click *Select*, then click *OK*.
- 9 When you are finished configuring the Subscriber object, click *OK* to exit the Subscriber object’s properties.

3.6.4 Updating Subscriber Configurations

The Subscriber software cannot run on a server if the Subscriber does not know its Tiered Electronic Distribution configuration, such as where it’s working directory is. Therefore, during the installation process, you determine a basic Tiered Electronic Distribution configuration for each of the Subscribers that you are installing.

Using this input, the installation program creates a `tednode.properties` file on each Subscriber server that contains the Subscriber’s initial Tiered Electronic Distribution configuration. Until a server receives its first Distribution, this `tednode.properties` file provides the server with its Tiered Electronic Distribution configuration information, so that it can function as a Subscriber.

A Subscriber server can only receive configuration information from a Distributor server whose Distributor object is in the same tree as the server’s Subscriber object. This is known as the trusted tree, which is established during the installation process. For information on knowing when the trusted tree is necessary, see [“Subscriber Software Configuration and Trusted Trees” on page 159.](#)

When a Distributor server sends a Distribution to a Subscriber server, the Distributor first checks to see if that Subscriber server has a current Tiered Electronic Distribution configuration in the form of a `tcpip.nlm` file. If this is the first time the Subscriber has received a Distribution, it does not have that file. The Distributor then sends the `tcpip.nlm` file to the Subscriber, and the `tednode.properties` file is no longer used by the Subscriber. Then the Distributor checks again to see if the Subscriber server has a current `tcpip.nlm` file. Upon confirmation from the Subscriber, the Distribution is sent. In other words, the Distributor never sends a Distribution to a Subscriber server whose configuration information is not current.

You can update the `tcpip.nlm` file any time you make configuration changes to the Subscriber object's properties. However, Subscribers do not read eDirectory, so when a change is made to the Subscriber, it must rely on the Distributor server to discover those changes and send the new configuration information to the Subscriber server, updating its `tcpip.nlm` file.

If you should install the Subscriber software to a server that does not have a Subscriber object in any eDirectory tree, such as a Microsoft domain server, the `tednode.properties` file is used by such servers, in lieu of having its Tiered Electronic Distribution configuration updated by a Distributor server. In this case, for configuration changes, you need to edit the server's `tednode.properties` file. For more information, see [“The Tednode.properties File Requirement” on page 161](#) and [Section 3.13, “Editing the Tednode.properties File,” on page 190](#).

3.6.5 Associating Subscribers with Channels

Before a Subscriber can receive a Distribution, you need to associate the Subscriber to the Channel holding the Distribution. You can do this either from the Subscriber or Channel object's properties:

- ♦ [“Associating a Channel with Multiple Subscribers” on page 154](#)
- ♦ [“Associating a Subscriber with Multiple Channels” on page 154](#)

Associating a Channel with Multiple Subscribers

To send a particular Distribution to many Subscriber servers:

- 1 In ConsoleOne, right-click the Channel object where the Distribution is listed, then click *Properties*.
- 2 Select the *Subscribers* tab, click *Add*, then add the needed Subscribers.
- 3 Select the *Schedule* tab and select a schedule.

The schedule determines when Distributions that have been received are extracted or installed.

For information on the available schedules, see [Chapter 8, “Scheduling,” on page 317](#).

- 4 Click *OK* to save the changes.

Associating a Subscriber with Multiple Channels

To subscribe a Subscriber server to multiple Channels for receiving different Distributions:

- 1 In ConsoleOne, right-click the Subscriber object, then click *Properties*.
- 2 Select the *Channels* tab, click *Add*, then add the needed Channels.
- 3 Select the *Schedule* tab and select a schedule.

The schedule determines when Distributions that have been received are extracted or installed.

For information on the available schedules, see [Chapter 8, “Scheduling,” on page 317](#).

- 4 Select the *Variables* tab, fill in the following fields, then click *OK*:

Variable name: Can be used to determine the location of the destination directory where the files are extracted. Enter the name of the variable exactly as you are using it within the %...% symbols.

Value: This is the value of the variable, which can be another variable's name.

Description: Text field to provide details about the variable.

For information on variables, see [Section 9.6, “Using Variables to Control File Extraction,”](#) on page 349.

5 Click *OK* to save the changes.

3.6.6 Deleting Subscriber Objects That Are Part of a Distributor’s Routing Hierarchy

If a Subscriber object is removed from eDirectory, or a Subscriber server is removed from the network (whether its Subscriber object is also removed or left in eDirectory), and that Subscriber was part of a Distributor’s routing hierarchy, you need to edit the Distributor object’s properties to adjust the routing hierarchy accordingly. Otherwise, Distributions that are sent through that parent Subscriber do not reach the designated Subscriber servers.

3.7 Subscriber Groups

The following sections provide concepts and instructions for the Subscriber Group object:

- ◆ [Section 3.7.1, “Understanding Subscriber Groups,”](#) on page 155
- ◆ [Section 3.7.2, “Creating and Configuring Subscriber Groups,”](#) on page 156

3.7.1 Understanding Subscriber Groups

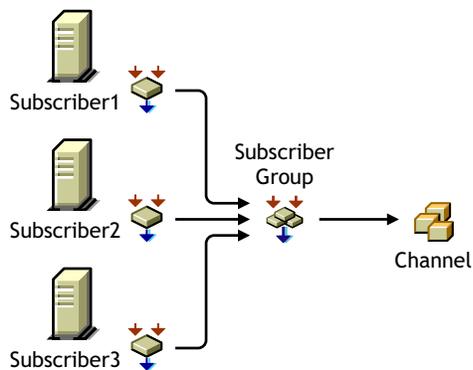
A Subscriber Group is an eDirectory object (TED Subscriber Group) used for grouping Subscribers objects.

- ◆ [“Functional Relationship with Other Tiered Electronic Distribution Objects”](#) on page 155
- ◆ [“Subscriber Group Description”](#) on page 155
- ◆ [“Scheduling”](#) on page 156

Functional Relationship with Other Tiered Electronic Distribution Objects

[Figure 3-30](#) illustrates a Subscriber Group’s relationship with Subscribers and Channels.

Figure 3-30 *Using a Subscriber Group*



Subscriber Group Description

A Subscriber Group is used for grouping Subscribers that have the same Distribution needs.

Subscriber Groups are useful when you are sending several different Distributions to the same set of Subscribers. There is no need to create a Subscriber Group if it is only associated with one Channel.

For example, Distribution A is in Channel A, Distribution B is in Channel B, and so on. Then, without using a Subscriber Group, you need to subscribe each of your Subscribers to Channel A, then each to Channel B, and so on, which could be a very long process. However, by using a Subscriber Group, you only need to create the group, add the Subscribers to it, then subscribe that one group to each Channel.

Another use of a Subscriber Group is that when the group is associated with two or more Channels, you can edit the group's membership more easily than making the same changes in multiple Channels. For example, to remove a Subscriber from one Subscriber Group, you just edit that one group's properties. To remove that same Subscriber from several Channels, you need to edit each Channel's properties.

Scheduling

Subscriber Groups are not scheduled.

3.7.2 Creating and Configuring Subscriber Groups

- 1 In ConsoleOne, select the container to hold the Subscriber Group object, click *File > New > Object*, then select *TED Subscriber Group*.
- 2 In the New TED Subscriber Group dialog box, provide a name for the Subscriber Group (worksheet [item 17](#)), select *Define additional properties*, then click *OK*.
- 3 Click *General > Settings* and provide a description.
- 4 To populate the group with Subscribers, select the *Members* tab and do the following:
 - 4a Click *Add*, browse for and select the Subscriber objects (worksheet [item 18](#)), then click *OK*.
 - 4b To remove any Subscribers from the list, select the Subscribers, then click *Delete*.
 - 4c To view the properties of any Subscriber, select the Subscriber, then click *Details*.
- 5 Click *OK* when you have finished configuring the Subscriber Group object.

3.8 External Subscribers

The following sections provide concepts and instructions for the External Subscriber object:

- ♦ [Section 3.8.1, "Understanding External Subscribers," on page 156](#)
- ♦ [Section 3.8.2, "Using External Subscribers for Out-of-Tree Distributions," on page 162](#)
- ♦ [Section 3.8.3, "Creating and Configuring External Subscribers," on page 164](#)

3.8.1 Understanding External Subscribers

An External Subscriber is an eDirectory object (TED External Subscriber) that represents a Subscriber object in another tree.

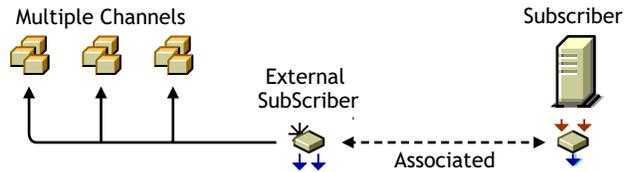
- ♦ ["Functional Relationship with Other Tiered Electronic Distribution Objects" on page 157](#)
- ♦ ["External Subscriber Description" on page 157](#)

- ◆ “Subscriber Software Configuration and Trusted Trees” on page 159
- ◆ “Scheduling” on page 162

Functional Relationship with Other Tiered Electronic Distribution Objects

Figure 3-31 illustrates an External Subscriber’s relationship with the Channel:

Figure 3-31 *External Subscriber’s Relationship with a Channel*



The External Subscriber subscribes to the Channels.

External Subscriber Description

A Distributor cannot send its Distributions to a Subscriber server whose Subscriber object is in a different tree than the Distributor’s object, or to a server that does not have a Subscriber object. An External Subscriber object is needed for out-of-tree distributions.

For information on the External Subscriber object, see the following:

- ◆ “The External Subscriber’s Purpose” on page 157
- ◆ “Distribution Information Not Maintained” on page 158
- ◆ “Duplicate Distribution Management” on page 158
- ◆ “External Subscriber Characteristics” on page 158
- ◆ “External Subscriber Requirements” on page 158
- ◆ “The External Subscriber Object’s Properties” on page 159

The External Subscriber’s Purpose

If you installed all of your Tiered Electronic Distribution objects in one tree, an External Subscriber object is not necessary, because you can send your Distributions using the Distributor and Subscriber objects that are in the same tree.

However, the External Subscriber object is useful for sending out-of-tree Distributions when one of the following conditions exists:

- ◆ **The target server has no Subscriber object in any tree:** The target server, such as a Windows server in a Microsoft domain, has only the Subscriber software installed on it.
- ◆ **The target server has a Subscriber object in a different tree:** The target server has the Subscriber software installed on it, but its Subscriber object is in a different tree than the Distributor object that is sending the Distribution.

Because the External Subscriber is only an object in an eDirectory tree, it does not actually handle the Distribution files; it simply identifies which server is to receive them.

Distribution Information Not Maintained

When sending any Distribution through External Subscribers, trusted tree rights cannot be maintained.

When sending Desktop Application Distributions through External Subscribers, application object associations cannot be maintained. However, it can send the group association, because it creates that.

Duplicate Distribution Management

You can use an External Subscriber object to circumvent the need to duplicate Distribution work in another tree.

For example, a few Subscribers on a tree at a remote site could receive all of their Distributions via the External Subscriber in the Distributor's tree. That would prevent the need to have a Distributor server at the remote site, including duplicating the Distribution configuration and management effort there.

External Subscriber Characteristics

An External Subscriber is associated with a server running the Subscriber software that has no Subscriber object in any tree, or no Subscriber object in the same eDirectory tree as the Distributor from which it receives the Distribution.

External Subscriber objects are associated with a Subscriber server through an IP address or DNS name of that server.

You can send Distributions outside of eDirectory, such as to a Windows server in a Microsoft domain. For more information on this type of Distribution, see [“Subscriber Software Configuration and Trusted Trees” on page 159](#) and [“The Tednode.properties File Requirement” on page 161](#).

External Subscriber objects cannot be parent Subscribers. If an External Subscriber has a parent Subscriber, both the External Subscriber's and parent Subscriber's objects must reside in the same tree.

External Subscriber Requirements

If a target server's Subscriber object is in a different tree from the Distributor object of the server that sends it a Distribution, that target server must be represented by an External Subscriber object in the Distributor's tree.

Because Tiered Electronic Distribution uses IP addresses or DNS names to locate servers, Subscriber objects can be in a different tree than those servers' NCP objects.

An External Subscriber must be subscribed to the Channel that lists the Distributions needed by its associated Subscriber.

The server receiving a Distribution via an External Subscriber must have the Subscriber software installed on it so it can receive and extract the Distribution. It is not required to have a Subscriber object in any tree, such as if it is a Windows server in a domain (see [“Subscriber Software Configuration and Trusted Trees” on page 159](#) and [“The Tednode.properties File Requirement” on page 161](#)).

The External Subscriber Object's Properties

The External Subscriber object properties contain only the following:

- ◆ IP address or DNS name of the Subscriber server in a different tree or a domain (required)
This is the ID of the Subscriber server in one tree that is to receive a Distribution from a Distributor in another tree (the tree where the External Subscriber object resides).
- ◆ The Channels it is subscribed to (required)
This is for identifying which Distributions need to be sent to the Subscriber server in the other tree.
- ◆ Membership in a Subscriber Group (optional)
You can use this for subscribing the External Subscriber to the Channels subscribed to by the group.
- ◆ Context of a parent Subscriber in the External Subscriber's own tree (optional)
A parent Subscriber is usually in the Distributor's distribution hierarchy.
If used, the parent Subscriber does the physical work in sending the Distribution file to the server in the other tree. Otherwise, the Distributor server sends the Distribution directly to the Subscriber server in the other tree.

Subscriber Software Configuration and Trusted Trees

Subscribers can be configured by a Distributor, but External Subscribers cannot. External Subscribers are just objects identifying a server. However, a Subscriber server identified by an External Subscriber object must have a Tiered Electronic Distribution configuration in order to receive the Distributions via the External Subscriber object.

Using the External Subscriber object brings up the need to understand trusted trees:

- ◆ [“The Reason for Trusted Trees” on page 159](#)
- ◆ [“Determining the Trusted Tree” on page 160](#)
- ◆ [“The Tednode.properties File Requirement” on page 161](#)
- ◆ [“Preventing Trusted Tree Errors for Policy Package Distributions” on page 161](#)

The Reason for Trusted Trees

The following applies to any NetWare or Windows server, whether it has an NCP server object in an eDirectory tree or a server object in a Microsoft domain:

- ◆ During installation, the server can have both a Subscriber object created for it and the Subscriber software installed to it
- ◆ During installation, the server can have only the Subscriber software installed to it (no Subscriber object is created)
- ◆ During installation, you should identify the trusted tree of any server that does not have a Subscriber object created for it

Identifying a trusted tree has two purposes:

- ◆ To locate a Distributor that can update the Subscriber's Tiered Electronic Distribution configuration information

- ◆ To indicate which tree to accept policies from

A Subscriber server's Tiered Electronic Distribution configuration information is stored in eDirectory in its Subscriber object (which the Distributor reads), and in a `tcpip.nlm` file in the Subscriber server's file system (which the Subscriber reads). A Distributor server sending the configuration information must have its Distributor object in the same tree as the Subscriber object that it is configuring.

A Subscriber server can receive its Subscriber software configuration only from a Distributor in its trusted tree. The trusted tree is where the server's Subscriber object and that Distributor object both reside. This is not the tree where an associated External Subscriber object resides, and it doesn't matter whether it's the same tree where the server's NCP object resides.

A Subscriber server that does not have a Subscriber object in any tree (such as a Windows server in a Microsoft domain), must use its `tednode.properties` file for its Tiered Electronic Distribution configuration information. This file is created on the server when you installed the Subscriber software. Then it can receive and extract Distributions from a Distributor in another tree (via an External Subscriber object). The extraction process is the time when the trusted tree requirement must be met. For more information, see [“The Tednode.properties File Requirement” on page 161](#).

Determining the Trusted Tree

There are two situations that deal with whether to install Subscriber objects for Subscriber servers:

- ◆ **eDirectory server:** When you install the Subscriber software to a server whose NCP server object is in another tree, you have one of the following options:
 - ◆ You can create the Subscriber object in the Distributor's tree, which might not be the tree where the Subscriber's NCP server object resides (the server's Subscriber and NCP objects do not need to be in the same tree). In this case, you do not need an External Subscriber object for sending Distributions to that Subscriber, because its object is not out-of-tree.

The Subscriber server's trusted tree is the same tree where the Distributor object resides. Therefore, it receives its Tiered Electronic Distribution configuration updates from the Distributor in its trusted tree.
 - ◆ You can elect to not create a Subscriber object for the server. In this case, you need to use the `tednode.properties` file to configure that Subscriber server. You also need to use an External Subscriber object to send Distributions to that server.

In order for this Subscriber to have policies enforced on it, you need to identify its trusted tree, which would be the tree it receives Policy Package Distributions from.
- ◆ **Non-eDirectory server:** When you install the Subscriber software to a server that is in a Microsoft domain, and therefore does not have an NCP server object in any eDirectory tree, you can create a Subscriber object for this server, but it is not required. If you do not have a Subscriber object, you need to use the `tednode.properties` file to configure that Subscriber server. You also need to use an External Subscriber object to send Distributions to this server.

In order for this Subscriber to have policies enforced on it, you need to identify its trusted tree, which would be the tree it receives Policy Package Distributions from.

The File Installation Paths and Options page in the installation program contains the Trusted Tree field. However, this field is only displayed if you deselect the Create eDirectory Objects check

box on the Installation Options page. This causes the installation program to install only software for the selected servers.

You must select a trusted tree for each server where you have selected to install the Subscriber software, or your Policy Package Distributions might not extract on that Subscriber server, because policies point to objects in a tree.

For installation instructions concerning the Trusted Tree field, see the steps in the applicable sections under “[Installation on NetWare and Windows Servers](#)” in the *Novell ZENworks 7 Server Management Installation Guide*.

The Tednode.properties File Requirement

A `tednode.properties` file must be used to provide configuration information for the following Subscriber servers:

- ◆ A Subscriber server that has a Subscriber object and has not yet received its first Distribution. After it does, it then uses the `tcpip.nlm` file given to it by the Distributor in its trusted tree that is sending that first Distribution, and it no longer uses the `tednode.properties` file.

A Subscriber can only be configured by a Distributor server whose object is in the same tree as the Subscriber’s object.

- ◆ A Subscriber server that does not have a Subscriber object in any tree.

This could be a Windows server in a Microsoft domain where you only installed the Subscriber software without creating the object.

If you installed the Subscriber software (using the ZENworks 7 Server Management installation program) without creating the Subscriber object, the `tednode.properties` file was automatically created and configured.

For more information, see [Section 3.13, “Editing the Tednode.properties File,” on page 190](#).

Preventing Trusted Tree Errors for Policy Package Distributions

In order to prevent trusted tree errors when sending a Policy Package Distribution to a Subscriber using an External Subscriber object, you must edit the `agentinfo.properties` file:

- 1 On the server using the External Subscriber object to receive a Policy Package Distribution, locate the `agentinfo.properties` file, which is in the `\zenworks\pds\ted` directory.
- 2 Open the `agentinfo.properties` file in a text editor.
- 3 Add the following lines in the file:

```
TRUSTED_TREE=source_tree_name
TRUSTED_TDN=External_Subscriber_DN
```

where `source_tree_name` is the tree where the Distributor object resides that is sending the Policy Package Distribution, and `External_Subscriber_DN` is the fully-distinguished name of the External Subscriber object receiving the Distribution.

- 4 Save the `agentinfo.properties` file and exit the text editor.
- 5 When ready, send the Policy Package Distribution to the External Subscriber.

Scheduling

The External Subscriber object is not scheduled.

3.8.2 Using External Subscribers for Out-of-Tree Distributions

Review the following sections to understand how to use External Subscribers for out-of-tree distributions:

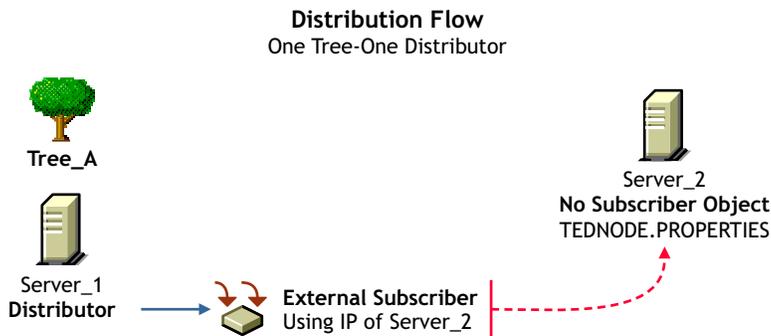
- ◆ “External Subscriber, One Distributor, and One Tree” on page 162
- ◆ “External Subscriber, Multiple Distributors, and Multiple Trees” on page 163

External Subscriber, One Distributor, and One Tree

After you install Policy and Distribution Services software to your servers, you can send Distributions to a server that does not have a Subscriber object in any tree using the External Subscriber object.

The Tiered Electronic Distribution configuration illustrated in [Figure 3-32](#) might exist for the Distributor’s routing of its Distributions through External Subscribers:

Figure 3-32 *Distribution Flow in One Tree*



In this example, Server_2 does not have a Subscriber object in any tree. It has only the Subscriber software installed on it so that it can receive and extract Distributions. It can be a NetWare server with an NCP server object in any tree, or a Windows server in a Microsoft domain.

To send a Distribution from Distributor_A to Server_2, create an External Subscriber object in Tree_A and list Server_2’s IP address or DNS name in the External Subscriber object’s properties.

- ◆ “The eDirectory Distribution View” on page 162
- ◆ “The Actual Distribution Process” on page 163
- ◆ “Configuring the Subscriber Server” on page 163
- ◆ “The Subscriber Server’s Trusted Tree” on page 163

The eDirectory Distribution View

From an eDirectory perspective, the Distribution is sent from the Distributor object to the External Subscriber object, which in turn sends it to Server_2. You can use a parent Subscriber in Tree_A (not shown) where you do not want the Distributor to be directly sending its Distributions to Server_2.

The Actual Distribution Process

From a topology perspective, the Distribution file is sent from Server_1 to Server_2, using the IP address or DNS name of Server_2 that is located in the External Subscriber object's properties.

Configuring the Subscriber Server

Server_2 receives its Tiered Electronic Distribution configuration information from the `tednode.properties` file installed on its server when the Subscriber software was installed there. Because there is no Subscriber object to configure, you need to edit Server_2's `tednode.properties` file in order to make configuration changes. For information on editing the `tednode.properties` file, see [Section 3.13, "Editing the Tednode.properties File," on page 190](#).

The Subscriber Server's Trusted Tree

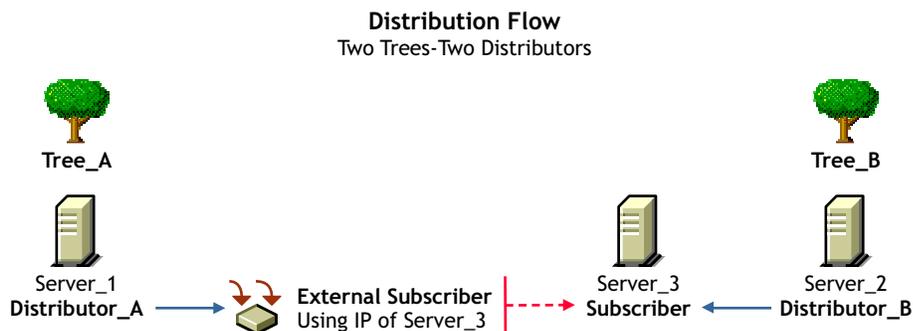
In order for Server_2 to have policies enforced on it, Tree_A needs to be established as its trusted tree during installation of the Subscriber software to the server. For the installation steps, see ["Installation on NetWare and Windows Servers"](#) in the *Novell ZENworks 7 Server Management Installation Guide*.

External Subscriber, Multiple Distributors, and Multiple Trees

After you install Policy and Distribution Services software to your servers in multiple trees, you can send Distributions between trees using the External Subscriber object.

The Tiered Electronic Distribution configuration illustrated in [Figure 3-33](#) might exist for the Distributor's routing of its Distributions through External Subscribers:

Figure 3-33 *Distribution Flow in Two Trees*



In this example, Server_3 has a Subscriber object in Tree_B.

To send a Distribution from Distributor_A to Server_3, you create an External Subscriber object in Tree_A and list Server_3's IP address or DNS name in the External Subscriber object's properties.

- ♦ ["The eDirectory Distribution View" on page 164](#)
- ♦ ["The Actual Distribution Process" on page 164](#)
- ♦ ["Subscriber Server_3's Trusted Tree and Its Tiered Electronic Distribution Configuration" on page 164](#)

The eDirectory Distribution View

From an eDirectory perspective, the Distribution is sent from Distributor_A to the External Subscriber object, which in turn sends it to Server_3. You can use a parent Subscriber in Tree_A (not shown) where you do not want Distributor_A to be directly sending its Distributions to Server_3.

The Actual Distribution Process

From a topology perspective, the Distribution file is sent from Server_1 to Server_3, using the IP address or DNS name of Server_3 that is located in the External Subscriber object's properties.

Subscriber Server_3's Trusted Tree and Its Tiered Electronic Distribution Configuration

Each tree has a Distributor that provides configuration information for the Subscriber servers in its own tree.

Server_3 receives its Tiered Electronic Distribution configuration information from Distributor_B, because Tree_B was set as Server_3's trusted tree when it was made a Subscriber using the installation program. However, Server_3 cannot extract a Distribution from Distributor_A until it has been configured by Distributor_B, which is done the first time the Subscriber receives a Distribution from Distributor_B.

3.8.3 Creating and Configuring External Subscribers

You can create External Subscriber objects for sending Distributions to Subscriber servers with Subscriber objects residing on other trees or to Subscriber servers that do not have a Subscriber object in any tree.

The following sections provide steps to create and configure an External Subscriber:

- ♦ [“Creating an External Subscriber Object” on page 164](#)
- ♦ [“Configuring the External Subscriber Object” on page 164](#)

Creating an External Subscriber Object

- 1 In ConsoleOne, select the container to hold the External Subscriber object, click *File > New > Object*, then select *TED External Subscriber*.
- 2 Provide a name for the External Subscriber object.
Make the name unique to help identify the server from the other tree.
- 3 Provide the server's TCP/IP address or DNS name, then click *OK*.
This must be a valid TCP/IP address or fully distinguished DNS name.
- 4 Continue with [“Configuring the External Subscriber Object” on page 164](#).

Configuring the External Subscriber Object

- 1 In ConsoleOne, right-click an External Subscriber object, then click *Properties*.
- 2 Click *General > Settings* and fill in the Setting fields:
Use policy: Select this check box if you want to use the values set in the Tiered Electronic Distribution policy that is being enforced on the External Subscriber's server.

If you select this option, the Parent Subscriber field is dimmed and the policy settings are used instead.

Parent Subscriber: Specifies a parent Subscriber from which all Distributions are received.

Because the routing hierarchy in a Distributor object's properties only accounts for parent Subscribers, this field is where you can connect an end-node Subscriber to the routing hierarchy. These end-node Subscribers (which in this case are External Subscribers) cannot be used to pass Distributions to other Subscribers.

3 Select the *Network Address* tab and verify the IP address of the External Subscriber's server.

IP address: You provided this IP address when you created the object. Verify that it is correct.

4 Select the *Channels* tab and fill in the fields:

♦ **Channels this Subscriber is subscribed to**

Lists the Channels the External Subscriber is subscribed to.

Active: To activate a Channel for this External Subscriber so it can receive the Channel's Distributions, select a Channel, then select the check box to enable it. To deactivate a Channel so that the External Subscriber does not receive the Channel's Distributions, deselect the check box to disable it.

Channel: Click Add to create a Channel. Click Details to edit a Channel.

♦ **Channels subscribed to through Subscriber Group memberships**

Lists the Subscriber Groups that the External Subscriber is a member of, paired with which Channels the External Subscriber is subscribed to by virtue of membership in a Subscriber Group.

These columns are for display only. The Details, Add, and Delete buttons do not apply.

Active: Indicates whether the Channel subscribed to is active.

Channel: Displays the Channel subscribed to through membership in a group.

Subscriber Groups: Displays the groups the External Subscriber is a member of. You can sort the listing by clicking the column heading.

5 To include this External Subscriber in a group, click *Group Membership*, click *Add*, browse for a Subscriber Group object, click *Select*, then click *OK*.

6 When you are finished configuring the External Subscriber object, click *OK* to exit the object's properties.

3.9 Configuring Multiple Tiered Electronic Distribution Objects

When you have the same configuration change to make to several Tiered Electronic Distribution objects, you can save time by modifying the properties of multiple objects.

You can perform multiple object modifications for the following Tiered Electronic Distribution objects:

| | | | |
|-------------|---------------------|--------------|----------------|
| Distributor | Subscriber Group | Distribution | Policy Package |
| Subscriber | External Subscriber | Channel | |

For more information, see:

- ♦ [Section 3.9.1, “Issues with Modifying Multiple Tiered Electronic Distribution Object Properties,” on page 166](#)
- ♦ [Section 3.9.2, “Modifying Multiple Tiered Electronic Distribution Object Properties,” on page 167](#)
- ♦ [Section 3.9.3, “Property Tabs Available for Multiple-Object Modifications,” on page 167](#)

3.9.1 Issues with Modifying Multiple Tiered Electronic Distribution Object Properties

- ♦ **Available properties:** Although the purpose is to provide a means to make the same changes to multiple objects, not all properties for the Tiered Electronic Distribution objects can be modified using this method.

The Schedule and Other property tabs are not available for editing the properties of multiple-selected Tiered Electronic Distribution objects. For the Distribution object, the Type tab is also not available. For changes to these property tabs, you must edit each Tiered Electronic Distribution object individually.

- ♦ **Modified fields:** The fields where you make changes in the Properties of Multiple Objects dialog box are the only modifications that are made for the selected objects. In other words, if you leave a field blank (you do not modify it), no change is made in that field for all of the selected objects. Each object retains its original field entry.

Where objects have different information in a given field, that field is blank in the Properties of Multiple Objects dialog box.

- ♦ **Removing information:** In some fields, a space is a valid entry. You can use this as a method for removing varied existing entries for each of the selected Tiered Electronic Distribution objects when you want the field to be blank for all of the selected objects.
- ♦ **Policy defaults:** If you have a Tiered Electronic Distribution policy in force, the Use Policy check box is displayed in each Tiered Electronic Distribution object’s properties, but only selected for the individual Tiered Electronic Distribution objects where the policy applies (because their properties have never been edited, or you selected that check box).

For multiple object properties, if the Use Policy check box is displayed and selected, the policy’s contents are displayed in dimmed text in the applicable fields. These attributes are only applicable to those Tiered Electronic Distribution objects whose individual properties contain a selected Use Policy check box.

You can deselect the Use Policy check box when editing multiple properties to disable the Tiered Electronic Distribution policy for the selected Tiered Electronic Distribution objects that were previously using the policy. Any changes you make are replicated to all selected Tiered Electronic Distribution objects and the Tiered Electronic Distribution policy are no longer in force for any of those objects.

IMPORTANT: If the Working Directory field for an object received its location from the Tiered Electronic Distribution policy, and you deselect the Use Policy check box when editing multiple properties, the Working Directory field is then left blank for that object. Therefore, the next time you access the properties for that object, you will be required to provide a working directory location.

3.9.2 Modifying Multiple Tiered Electronic Distribution Object Properties

To modify the properties of multiple Tiered Electronic Distribution objects:

- 1 In ConsoleOne, select a number of Tiered Electronic Distribution objects.

They must be of the same type, such as all Distributor objects. The Properties of Multiple Objects menu option do not display if you select multiple objects of different types.

You can select multiple objects using the Shift and Ctrl keys.

- 2 Right-click the selected objects and click *Properties of Multiple Objects*.

Each of the selected objects is listed in the Objects to Modify tab on the Properties of Multiple Objects dialog box. These are the objects that have their properties modified when you make changes.

- 3 To change the objects displayed in the list, click *Add* or *Remove*.

The Add button allows you to browse for other Tiered Electronic Distribution objects. Only objects of the type you have previously selected are displayed for adding to the list.

Before selecting the Remove button, you must first select one or more objects in the list. This only removes the objects from the list, not from eDirectory.

- 4 Select a tab containing the property that you want to modify.

For descriptions of the property tabs available for the various Tiered Electronic Distribution objects, see [Section 3.9.3, “Property Tabs Available for Multiple-Object Modifications,” on page 167](#).

- 5 Edit the property.

The changes are made to all of the objects listed in the *Objects to Modify* tab.

For more information on individual property fields, see the descriptions within the steps in the following sections:

- ♦ [“Configuring Distributors” on page 106](#)
- ♦ [“Creating a Distribution” on page 123](#)
- ♦ [“Creating and Configuring Channels” on page 146](#)
- ♦ [“Configuring Subscribers” on page 150](#)
- ♦ [“Creating and Configuring Subscriber Groups” on page 156](#)
- ♦ [“Creating and Configuring External Subscribers” on page 164](#)

- 6 Repeat [Step 4](#) and [Step 5](#) until you have finished modifying the various properties for the selected objects.

- 7 When finished modifying properties, click *OK* to close the Properties of Multiple Objects dialog box.

All changes that you have made are updated for all of the selected objects.

3.9.3 Property Tabs Available for Multiple-Object Modifications

The tables in the following sections list the property tabs that are available in the multiple object editing mode for each Tiered Electronic Distribution object.

IMPORTANT: Generally, if you change information, it is changed for all of the selected objects. Exceptions are noted in the explanations.

- ◆ “Distributor Object” on page 168
- ◆ “Distribution Object” on page 168
- ◆ “Channel Object” on page 169
- ◆ “Subscriber Object” on page 170
- ◆ “External Subscriber Object” on page 170
- ◆ “Subscriber Group Object” on page 171
- ◆ “Policy Package Object” on page 172

Distributor Object

Table 3-9 *Distributor Object Property Tabs*

| Property Tabs Available | Explanation |
|-------------------------|---|
| Objects to Modify | You can add or remove Distributor objects from the list of objects to be modified. This does not add or remove the objects from eDirectory. |
| General | This includes the Settings and Messaging subtabs. For the Settings subtab, none of the fields display information, even if it is identical between the selected Subscriber objects. However, dimmed text is displayed in fields where the Tiered Electronic Distribution policy is in effect for one or more of the selected Tiered Electronic Distribution objects. In the Settings subtab, you can only add new information that applies to all of the selected Subscriber objects. In the Messaging subtab, you can edit existing entries. |
| Routing | If there are any differences in routing hierarchies between the selected Distributor objects, nothing is displayed for this tab. You can only edit routing hierarchies for multiple Distributor objects when they are identical. |
| NDS Rights | This tab includes the New Trustees and the Inherited Filter Rights subtabs. |

Distribution Object

Table 3-10 *Distribution Object Property Tabs*

| Property Tabs Available | Explanation |
|-------------------------|--|
| Objects to Modify | You can add or remove Distribution objects from the list of objects to be modified. This does not add or remove the objects from eDirectory. |

| Property Tabs Available | Explanation |
|-------------------------|--|
| General | <p>This includes the Settings and Restrictions subtabs.</p> <p>For the Settings subtab, none of the fields display information, even if it is identical between the selected Subscriber objects. However, dimmed text is displayed in fields where the Tiered Electronic Distribution policy is in effect for one or more of the selected Tiered Electronic Distribution objects.</p> <p>In the Settings subtab, you can only add new information that applies to all of the selected Subscriber objects. In the Restrictions subtab, you can edit existing entries.</p> |
| Channels | <p>Channels do not automatically display on this tab. You can only browse for Channels to add to each of the selected Distribution objects, or browse for a Channel to be removed from each of the selected Distribution objects that are associated with that Channel.</p> <p>Adding or removing a Channel in the list on this tab does not add or remove the Channel object from eDirectory.</p> |
| NDS Rights | This tab includes the New Trustees and the Inherited Filter Rights subtabs. |

Channel Object

Table 3-11 *Channel Object Property Tabs*

| Property Tabs Available | Explanation |
|-------------------------|---|
| Objects to Modify | You can add or remove Channel objects from the list of objects to be modified. This does not add or remove the objects from eDirectory. |
| General | <p>This includes the Settings subtab (with the Active check box and the Description field).</p> <p>For the Settings subtab, none of the fields display information, even if it is identical between the selected Subscriber objects. However, dimmed text is displayed in fields where the Tiered Electronic Distribution policy is in effect for one or more of the selected Tiered Electronic Distribution objects.</p> <p>In the Settings subtab, you can only add new information that applies to all of the selected Subscriber objects.</p> |
| Distributions | <p>Distributions do not automatically display on this tab. You can only browse for Distributions to add to each of the selected Channel objects, or browse for a Distribution to be removed from each of the selected Channel objects that are associated with that Distribution.</p> <p>Adding or removing a Distribution in the list on this tab does not add or remove the Distribution object from eDirectory.</p> |
| Subscribers | <p>Subscribers do not automatically display on this tab. You can only browse for Subscribers to add to each of the selected Channel objects, or browse for a Subscriber to be removed from each of the selected Channel objects that are associated with that Subscriber.</p> <p>Adding or removing a Subscriber in the list on this tab does not add or remove the Subscriber object from eDirectory.</p> |

| Property Tabs Available | Explanation |
|-------------------------|---|
| NDS Rights | This tab includes the New Trustees and the Inherited Filter Rights subtabs. |

Subscriber Object

Table 3-12 *Subscriber Object Property Tabs*

| Property Tabs Available | Explanation |
|-------------------------|--|
| Objects to Modify | You can add or remove Subscriber objects from the list of objects to be modified. This does not add or remove the objects from eDirectory. |
| General | <p>This includes the Settings, Messaging, and Working Context subtabs.</p> <p>For the Settings subtab, none of the fields display information, even if it is identical between the selected Subscriber objects. However, dimmed text is displayed in fields where the Tiered Electronic Distribution policy is in effect for one or more of the selected Tiered Electronic Distribution objects.</p> <p>In the Settings subtab, you can only add new information that applies to all of the selected Subscriber objects. In the Messaging subtab, you can edit existing entries.</p> |
| Channels | <p>Channels do not automatically display on this tab. You can only browse for Channels to add to each of the selected Subscriber objects, or browse for a Channel to be removed from each of the selected Subscriber objects that are associated with that Channel.</p> <p>Adding or removing a Channel in the list on this tab does not add or remove the Channel object from eDirectory.</p> |
| Variables | You can only add a new variable for all of the selected objects. Variables that are common among all of the selected objects are not displayed for editing. You must visit each Subscriber object individually to modify existing variables. |
| Group Membership | <p>Group Memberships do not automatically display on this tab. You can only browse for Group Memberships to add to each of the selected Subscriber objects, or browse for a Group Membership to be removed from each of the selected Subscriber objects that are associated with that Group Membership.</p> <p>Adding or removing a Group Membership in the list on this tab does not add or remove the Group Membership object from eDirectory.</p> |
| NDS Rights | This tab includes the New Trustees and the Inherited Filter Rights subtabs. |

External Subscriber Object

Table 3-13 *External Subscriber Object Property Tabs*

| Property Tabs Available | Explanation |
|-------------------------|---|
| Objects to Modify | You can add or remove External Subscriber objects from the list of objects to be modified. This does not add or remove the objects from eDirectory. |

| Property Tabs Available | Explanation |
|-------------------------|--|
| General | <p>This includes the Settings subtab.</p> <p>For the Settings subtab, only the Parent Subscriber field exists. If you make an entry here, all selected External Subscribers will have the same parent Subscriber.</p> |
| Channels | <p>Channels do not automatically display on this tab. You can only browse for Channels to add to each of the selected External Subscriber objects, or browse for a Channel to be removed from each of the selected External Subscriber objects that are associated with that Channel.</p> <p>Adding or removing a Channel in the list on this tab does not add or remove the Channel object from eDirectory.</p> |
| Group Membership | <p>Group Memberships do not automatically display on this tab. You can only browse for Group Memberships to add to each of the selected Subscriber objects, or browse for a Group Membership to be removed from each of the selected Subscriber objects that are associated with that Group Membership.</p> <p>Adding or removing a Group Membership in the list on this tab does not add or remove the Group Membership object from eDirectory.</p> |
| NDS Rights | This tab includes the New Trustees and the Inherited Filter Rights subtabs. |

Subscriber Group Object

Table 3-14 *Subscriber Group Object Property Tabs*

| Property Tabs Available | Explanation |
|-------------------------|---|
| Objects to Modify | You can add or remove Subscriber Group objects from the list of objects to be modified. This does not add or remove the objects from eDirectory. |
| General | This includes the Settings and Messaging subtabs. |
| Channels | <p>Channels do not automatically display on this tab. You can only browse for Channels to add to each of the selected Subscriber Group objects, or browse for a Channel to be removed from each of the selected Subscriber Group objects that are associated with that Channel.</p> <p>Adding or removing a Channel in the list on this tab does not add or remove the Channel object from eDirectory.</p> |
| Group Members | <p>Group Members do not automatically display on this tab. You can only browse for Group Members to add to each of the selected Subscriber objects, or browse for Group Members to be removed from each of the selected Subscriber objects that are associated with that Group Membership.</p> <p>Adding or removing a Group Membership in the list on this tab does not add or remove the Group Membership object from eDirectory.</p> |
| NDS Rights | This tab includes the New Trustees and the Inherited Filter Rights subtabs. |

Policy Package Object

Table 3-15 Policy Package Object Property Tabs

| Property Tabs Available | Explanation |
|-------------------------|---|
| Objects to Modify | You can add or remove Policy Package objects from the list of objects to be modified. This does not add or remove the objects from eDirectory. |
| Policies | This includes the various supported platform subtabs. For more information on the policies available on these platforms, see Section 4.1.6, “Server Policy Descriptions,” on page 200. |
| Distributions | Distributions do not automatically display on this tab. You can only browse for Distributions to add to each of the selected Policy Package objects, or browse for a Distribution to be removed from each of the selected Policy Package objects that are associated with that Distribution. Adding or removing a Distribution in the list on this tab does not add or remove the Distribution object from eDirectory. |
| NDS Rights | This tab includes the New Trustees and the Inherited Filter Rights subtabs. |

3.10 Sending Distributions

For information on sending Distributions, see the following:

- ◆ [Section 3.10.1, “Understanding the Distribution Processes,”](#) on page 172
- ◆ [Section 3.10.2, “Forcing a Single Distribution To Be Sent,”](#) on page 173
- ◆ [Section 3.10.3, “Sending Distributions Through Parent Subscribers,”](#) on page 174
- ◆ [Section 3.10.4, “Sending Distributions between Trees,”](#) on page 174
- ◆ [Section 3.10.5, “Sending Distributions: Firewall and Cluster Issues,”](#) on page 175

3.10.1 Understanding the Distribution Processes

Following are the processes for creating and sending a Distribution, generally done in this order:

- 1. Configure and schedule the Distributors.** You must use the installation program on the *ZENworks 7 Server Management with Support Pack 1 Program CD* to create a Distributor.

For information on Distributors, see [Section 3.3, “Distributors,”](#) on page 95 and [“Distributor Object’s Refresh Schedule”](#) on page 321.

- 2. Configure and schedule the Subscribers.** You must use the installation program on the *ZENworks 7 Server Management with Support Pack 1 Program CD* to create a Subscriber.

One of the primary configurations that you must do for Subscribers is to associate them with the Channels that hold the Distributions they need. For more information, see [Section 3.6.5, “Associating Subscribers with Channels,”](#) on page 154.

For information on Subscribers, see [Section 3.6, “Subscribers,”](#) on page 147 and [“Subscriber Object’s Extract Schedule”](#) on page 324.

- 3. Configure the necessary policies.** Policy Packages that contain the desired policies must be created in ConsoleOne or iManager before they are distributed.

For information on policies, see [Section 4.3, “Configuring Server Policies,” on page 205](#).

- 4. Create, configure, and schedule the Distributions.** You can use either ConsoleOne or iManager to create Distribution objects.

This could be the most time-consuming portion of the whole process, depending on the complexity of the Distribution to be configured. After you set up your Distributors and Subscribers and create the Distribution objects, you only need to utilize the Distributors’ routing hierarchies for distributing the files and policies to your Subscriber servers.

The Distribution object’s schedule is the best place to prevent an individual Distribution from being sent.

For information on Distributions, see [Section 3.4, “Distributions,” on page 110](#) and [“Distribution Object’s Build Schedule” on page 322](#).

- 5. Create, configure, and schedule the Channels.** You can use either ConsoleOne or iManager to create Channel objects.

Usually, you create a new Channel for each Distribution. It is generally easier to manage your distribution system by matching Channels with what they distribute. However, you can include multiple Distributions in a Channel, such as when they are related and all Subscribers subscribing to the Channel need all of those Distributions. For example, a Channel could hold several Distributions that each contain a different virus pattern update.

The Channel object is normally the best object to use for controlling whether Distributions should be sent. Setting its schedule to Never effectively stops the distribution process for all of the Distributions listed in it.

For information on Channels, see [Section 3.5, “Channels,” on page 144](#) and [“Channel Object’s Send Schedule” on page 323](#).

The Distributions are built, sent, and extracted according to the schedules that you set for each of the Tiered Electronic Distribution objects involved.

For information on the distribution processes, see [Section 3.2.2, “The Basic Distribution Process,” on page 88](#).

You might have accomplished some of the above processes during installation of Server Management and during your initial system configuration (see [Chapter 1, “Post-Installation Setup,” on page 29](#)).

3.10.2 Forcing a Single Distribution To Be Sent

If you want to send a single Distribution outside of the normal Refresh, Build, and Send schedules, and the Channel’s Send schedule is not ready to fire, you can manually force this distribution process using only the ZENworks Server Management role in iManager.

To force a single Distribution to be sent, do one of the following:

- ◆ If the Send Distribution Immediately After Building option is selected in the Distribution’s properties, go to iManager, click Distribution, then click Build Distribution.

Even if there are other Distributions in the Channel where this Distribution is listed, only this Distribution is sent.

- ◆ If the Send Distribution Immediately After Building option is not selected in the Distribution’s properties, go to iManager, click Distribution, click Build Distribution, click Channel, then click Distribute Channel.

All other Distributions in the Channel are also be sent if needed by the Subscribers.

As soon as a Subscriber receives an entire Distribution, it extracts it according to the Subscriber's Extract schedule.

3.10.3 Sending Distributions Through Parent Subscribers

Subscribers can receive and extract Distributions, and they can also pass on Distributions to other Subscribers. Subscribers that pass on Distributions are known as parent Subscribers.

Parent Subscribers do not need to be subscribed to the Distributions they are passing on. They simply receive a Distribution for passing it on to a subordinate Subscriber that has done two things:

- ◆ Subscribed to the Channel listing the Distribution
- ◆ Identified the parent Subscriber in the subordinate Subscriber's object properties

To set up parent Subscribers for passing on Distributions:

- 1 Determine a Subscriber object (hereafter referred to as "child Subscriber") that cannot receive a certain Distribution because this child Subscriber is not contained in the Distributor's routing hierarchy (the Distributor owning this Distribution).
- 2 In that Subscriber object's properties, click *General > Settings*, in the *Parent Subscriber* field browse for and select a Subscriber object that is contained in the Distributor's routing hierarchy, then click *OK*.

This establishes the Subscriber selected as a parent Subscriber. This distinction is not kept in the parent Subscriber's object properties, but only in the child Subscriber's.

- 3 Create a *Channel* object where only the child Subscriber is associated.
- 4 Create a Distribution, then associate it with the child Subscriber's Channel.
- 5 Send this Distribution.

Because this Distribution is associated only with the Channel where the child Subscriber is subscribed, the parent Subscriber does not extract it, but only passes it on to the child Subscriber.

Because the parent Subscriber is in the routing hierarchy of the Distributor, it has access to the Distribution for passing it on. However, the child Subscriber does not have any access to the Distributor, so it needs the parent Subscriber to provide access to the Distribution.

Although you can establish a parent Subscriber for a child Subscriber, the child Subscriber can still be subscribed to a Channel where the parent Subscriber is subscribed. Both Subscribers can receive and extract that Channel's Distributions without the parent Subscriber passing it on to the child Subscriber, because the child can have access to that particular Distributor's routing hierarchy. The key is whether the Distributor owning the desired Distribution can send it to the child Subscriber without using a parent Subscriber.

3.10.4 Sending Distributions between Trees

Using External Subscribers, you can send Distributions from one tree to another. To accomplish this, do the following:

- 1 Make sure Tiered Electronic Distribution is installed to both trees.

In the remaining steps, TREE1 represents the tree where the Distribution is created and TREE2 represents the other tree where you want the Distribution sent.

The server in TREE2 that is to receive the Distribution from TREE1 must have the Subscriber software installed on it (meaning it is a Subscriber in TREE2).

For information on installing Tiered Electronic Distribution, see “[Installation on NetWare and Windows Servers](#)” in the *Novell ZENworks 7 Server Management Installation Guide*.

2 In TREE1, create an External Subscriber object.

Make sure that the IP address or DNS name you provide for this object matches the Subscriber server in TREE2 where you want the Distribution to be sent.

For steps in creating External Subscribers, see [Section 3.8.3, “Creating and Configuring External Subscribers,”](#) on page 164.

3 In TREE1, create the Channel for the Distribution.

For steps in creating Channels, see [Section 3.5.2, “Creating and Configuring Channels,”](#) on page 146.

4 Associate the External Subscriber object you created in [Step 2](#) with the Channel you created in [step Step 3](#).

Other Subscribers from TREE1 can already be associated with this Channel.

For steps in associating Subscribers with Channels, see [Section 3.6.5, “Associating Subscribers with Channels,”](#) on page 154.

5 In TREE1, create the Distribution.

For steps in creating Distributions, see [Section 3.4, “Distributions,”](#) on page 110.

6 Associate this Distribution with the Channel you created in [Step 3](#).

7 Verify that the External Subscriber server in TREE2 received the Distribution.

3.10.5 Sending Distributions: Firewall and Cluster Issues

To send Distributions across a firewall, you must enable both the primary and secondary IP addresses of the servers running the Site List server or Distributor server software. If you only allow the secondary IP address to pass through the firewall, the Distribution cannot be sent because Tiered Electronic Distribution uses the primary IP addresses of its recipient servers.

If you are running ZENworks in a cluster, you also need to allow access to all primary IP addresses of all nodes involved.

3.11 Miscellaneous Tiered Electronic Distribution Issues

- ◆ [Section 3.11.1, “Directory Sync Granularity for File Distributions,”](#) on page 176
- ◆ [Section 3.11.2, “Understanding Dependencies in Tiered Electronic Distribution,”](#) on page 181
- ◆ [Section 3.11.3, “System Resources and Server Behavior,”](#) on page 181
- ◆ [Section 3.11.4, “Controlling I/O Rates and Concurrent Distributions,”](#) on page 182
- ◆ [Section 3.11.5, “Minimizing Messaging Traffic,”](#) on page 183

- ◆ Section 3.11.6, “Changing DNS Names or IP Addresses for Tiered Electronic Distribution Servers,” on page 185
- ◆ Section 3.11.7, “When a Tiered Electronic Distribution Process Fails,” on page 185

3.11.1 Directory Sync Granularity for File Distributions

The File Distribution has been enhanced with directory sync granularity:

- ◆ “Understanding Synchronization and Directory Sync Granularity” on page 176
- ◆ “How the Synchronization and Directory Sync Granularity Processes Work” on page 176
- ◆ “Synchronizing Directories for a File Distribution” on page 179

Understanding Synchronization and Directory Sync Granularity

A File Distribution, with or without synchronization enabled, adds or updates files and directories on a Subscriber server. However, with synchronization enabled it also causes the deletion of files and directories. Therefore, file and directory deletion on the Subscriber server is the main function of synchronization.

Directory sync granularity allows you to specify synchronization at any directory level in the Distribution to provide synchronization “from here down.”

How the Synchronization and Directory Sync Granularity Processes Work

Table 3-16 illustrates what a synchronized File Distribution does to the Subscriber server’s file system if synchronization is enabled in the Distribution:

Table 3-16 *Directory Sync Granularity Comparison*

| Files and directories located on the Distributor server that are contained in the File Distribution: | Applicable directories on the Subscriber server before the Distribution is received and extracted: |
|---|---|
| data:\zenworks\viruspatterns data:\zenworks\nw65sp\nw65sp1.exe data:\zenworks\nw65sp\nw65sp2 | data:\zenworks\viruspatterns data:\zenworks\nw65sp |
| Each of the end items in the above paths are synchronized in this Distribution. | One of the files is missing from the \viruspatterns directory on the Subscriber, and it is also missing the \nw65sp2 directory. |

Upon extraction of the File Distribution, the following occurs on the Subscriber server’s file system:

1. The missing virus pattern file is restored in the \viruspatterns directory.
The \viruspatterns directory is also made to exactly match the files and subdirectories contained in the Distribution by deleting any files or subdirectories on the Subscriber that are not contained in the Distribution.
2. The nw65sp1.exe file is updated in the \nw65sp directory. Nothing else is synchronized in that directory, because synchronization was not enabled for the \nw65sp directory itself.
3. The \nw65sp2 directory and its files and subdirectories are restored from the Distribution under the \nw65sp directory on the Subscriber.

Directory sync granularity also allows you to retain unsynchronized directories while synchronizing their peer directories. For example, you could select to synchronize the `data:\zenworks\viruspatterns` and `data:\zenworks\nw65sp\nw65sp2` directories, but not the `data:\zenworks\nw65sp` directory.

However, if you synchronize the parent `data:\zenworks` directory, all of its subdirectories must also be synchronized, because synchronization occurs from the specified directory and downward. Therefore, when you select directories to be synchronized, you cannot select a parent directory to be synchronized, then select some of its child directories to not be synchronized.

All child directories are automatically synchronized when a parent directory is set to be synchronized, and a parent directory automatically loses its synchronization when one of its child directories has synchronization turned off for it.

For example, [Table 3-17](#) illustrates this:

Table 3-17 Directory Sync Granularity Plan Comparisons

| Incorrect Plan | Correct Plan |
|--|---|
| Synchronize: <code>data:\zenworks</code> | Synchronize: <code>data:\zenworks\viruspatterns</code> <code>data:\zenworks\nw65sp\nw65sp1.exe</code> <code>data:\zenworks\nw65sp\nw65sp2</code> |
| Do not synchronize: <code>data:\zenworks\nw65sp</code> | Do not synchronize: <code>data:\zenworks</code> <code>data:\zenworks\nw65sp</code> |
| This does not work, because by synchronizing the <code>\zenworks</code> directory, the <code>\nw65sp</code> directory must also be synchronized. | This works, because the directories desired to not be synchronized are higher in the path than those that are desired to be synchronized. |

The next few sections describe synchronization scenarios:

- ◆ [“Synchronizing All Directories Under the Distribution’s Target Directory” on page 177](#)
- ◆ [“Using Directory Sync Granularity to Synchronize Directories at Various Levels” on page 178](#)
- ◆ [“Synchronizing a Subscriber Server Directory with Certain Distributor Server Files” on page 178](#)

Synchronizing All Directories Under the Distribution’s Target Directory

For the source, choose the directories on the Distributor server’s file system to be synchronized. For the destination, the directories to be synchronized might or might not already exist in the Subscriber server’s file system.

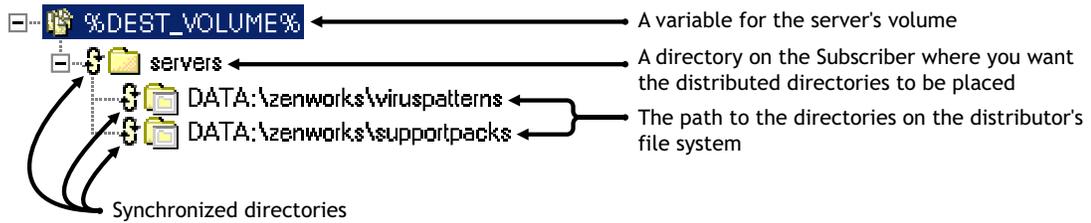
If the files already exist, when the Distribution is sent, their content is made to match that of the corresponding directory on the Distributor server’s file system. If they do not exist, they are added on the Subscriber server’s file system.

Determine where these directories should exist on the Subscriber server’s file system. In other words, there is a parent directory under which the synchronized directories are located, or there are

the synchronized directories located at the root of the Subscriber's file system. Variables can be used to specify the target on the Subscriber server's file system where the Distribution is to be extracted.

In **Figure 3-34**, the `\viruspatterns` and `\supportpacks` directories on the Distributor server are created and synchronized under the `vol1:\servers` directory on the Subscriber server.

Figure 3-34 *Synchronizing All Directories*

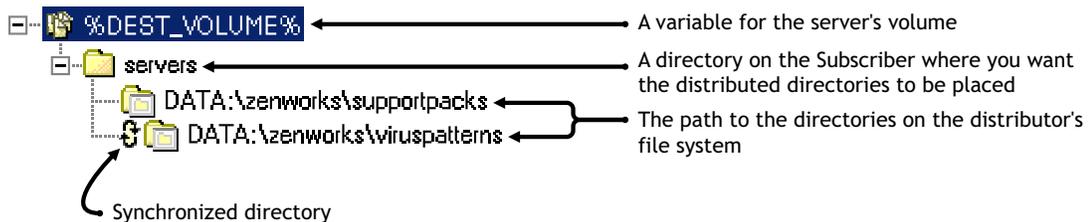


Using Directory Sync Granularity to Synchronize Directories at Various Levels

You can synchronize at the target level (`%DEST_VOLUME%`). In the example above, we have defined it as `vol1:\servers`. In that case, the only subdirectories that will exist under that directory are `viruspatterns` and `supportpacks`. All other existing subdirectories are deleted.

To retain other directories under `vol1:\servers`, you would not enable synchronization at the target level (`%DEST_VOLUME%`). Instead, you would drop down to the subordinate directories and synchronize those. For example, **Figure 3-35** illustrates synchronizing only the `\viruspatterns` directory. That means there could be other directories under `vol1:\servers` that would be unsynchronized, such as `\supportpacks`.

Figure 3-35 *Using Directory Sync Granularity*

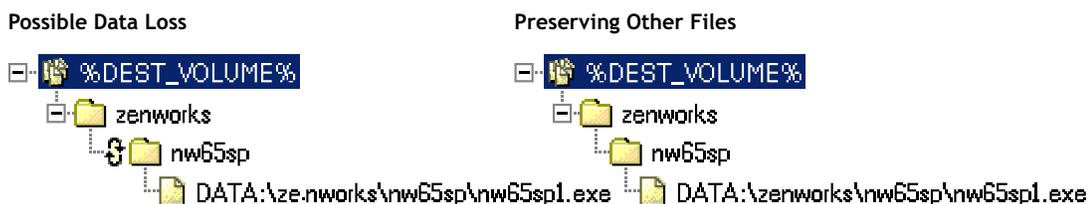


Synchronizing a Subscriber Server Directory with Certain Distributor Server Files

If the File Distribution has certain files on the Distributor server selected to be associated with a directory in the Distribution, you would not normally synchronize that directory in the Distribution. If you did, you could lose valuable data in that directory on the Subscriber server.

For example, **Figure 3-36** illustrates this:

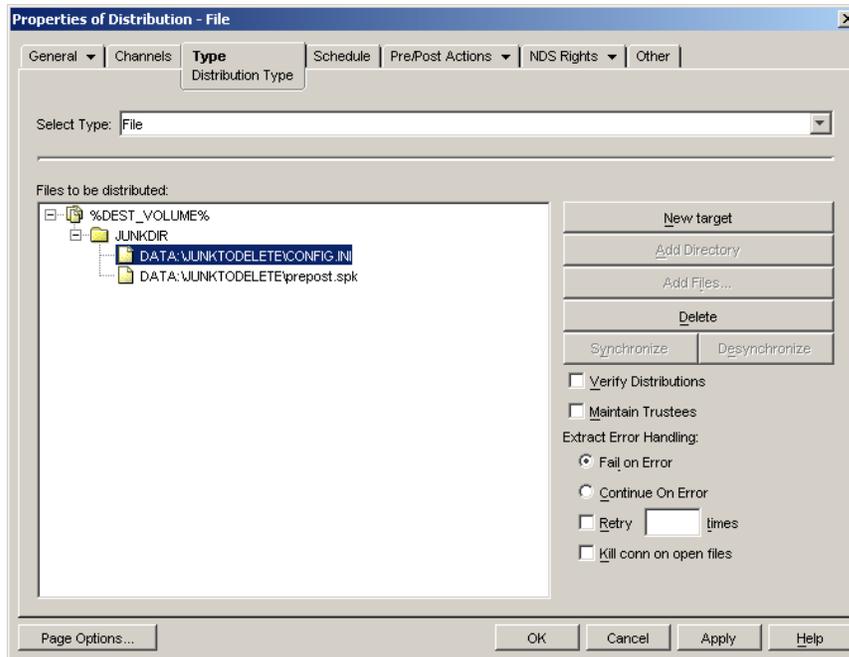
Figure 3-36 *Synchronizing a Subscriber Server Directory*



The Possible Data Loss method would cause all files in the \nw65sp directory and any of its subdirectories to be deleted from the Subscriber server, except for the nw65sp1.exe file. The Preserving Other Files method just updates the nw65sp1.exe file in the \nw65sp directory, leaving all other files and subdirectories unchanged on the Subscriber server.

Synchronizing Directories for a File Distribution

- 1 In ConsoleOne®, right-click a Distribution object, then click *Properties*.
- 2 Select the *Type* tab, then select *File* for the type of Distribution.



- 3 Click *New Target* and %DEST_VOLUME% as the default variable that contains the target path on each Subscriber server.

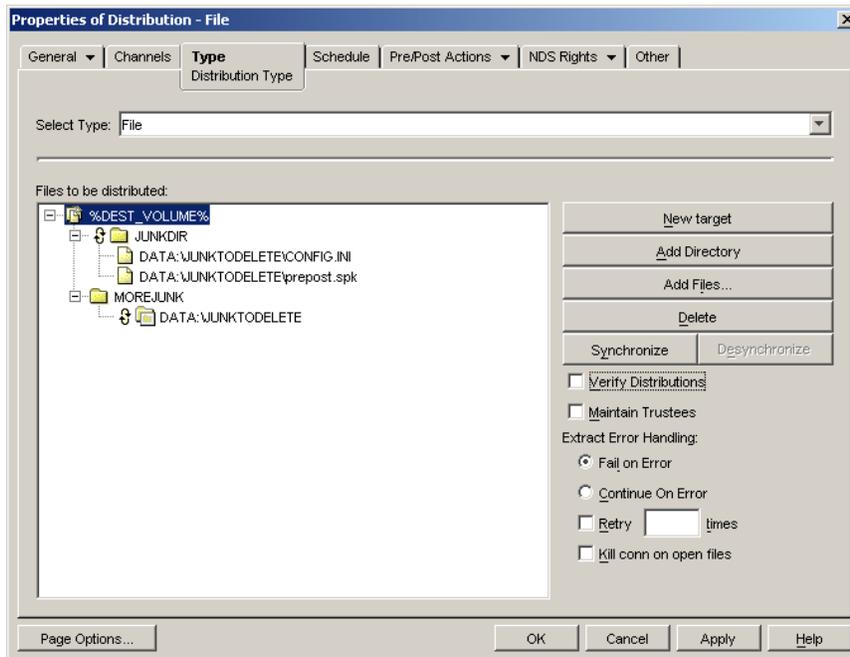
WARNING: If you want to synchronize at the target level, make sure this variable does not contain the root of the Subscriber server's file system.

- 4 Select the *Add Directory* and *Add Files* buttons to create the directory structure in the *Files To Be Distributed* box.

The *Add Directory* button is for the Subscriber server's target paths and directories, and the *Add Files* button is for browsing the distributor's file system and selecting directories and files that are to be included in the Distribution.

- 5 If necessary, click the plus signs to expand the directory structure.

By default, no directories in the listing are selected for synchronization.



- 6 To specify a directory to synchronize, select the directory in the *Files To Be Distributed* box, then click the *Synchronize* button.

The *Synchronize* icon (⊕) is placed in front of that directory name.

The *Synchronize* and *Desynchronize* buttons are dimmed when you select a filename instead of a directory name. You can only synchronize directories.

Because directory synchronization is done for the directory you select and all of its subdirectories (in other words, “from here down”), there is no need to synchronize any directories below a directory that you select for synchronization.

When you synchronize a directory and expand the directories below it, the *Synchronize* icon is displayed before each directory’s name.

- 7 To reverse your selection of a directory to be synchronized, select the directory name that has a *Synchronize* icon, then click the *Desynchronize* button.

If you desynchronize a directory and its parent directory was synchronized, the parent directory is also automatically desynchronized. This includes all directories (grandparents) up the path that might have been synchronized. In other words, you cannot synchronize a directory, then desynchronize any directory below it without also causing the directory to be desynchronized.

- 8 Repeat [Step 6](#) for each directory to synchronize.
- 9 Continue with configuring the Distribution.

For more information on configuring Distributions, see [“Creating and Configuring the Distribution” on page 57](#).

- 10 Click *OK* when finished configuring the Distribution.

For this Distribution, directory synchronization occurs with only the directories where you placed the *Synchronize* icon (⊕), including all of their subdirectories.

3.11.2 Understanding Dependencies in Tiered Electronic Distribution

Policy and Distribution Services agents (Policy/Package Agent and Distributor Agent) are dependent on one another and upon eDirectory. It is important to understand the following dependencies when using Policy and Distribution Services to manage your network:

- ◆ “Synchronization of Tiered Electronic Distribution Objects in eDirectory” on page 181
- ◆ “Unloading Parent Subscribers” on page 181

Synchronization of Tiered Electronic Distribution Objects in eDirectory

Server Management uses eDirectory as the repository for information needed by the Tiered Electronic Distribution and Server Policies components. Because eDirectory is a distributed database and can have partitions and replicas throughout the network, it takes time to synchronize all of the replicas each time Server Management objects are created or modified.

Unloading Parent Subscribers

You must change the parent Subscriber attribute in the Subscriber object to change the parent Subscriber. Then, the next time a Distribution is sent, the distribution route to the Subscriber reflects the new parent Subscriber.

If a parent Subscriber Java process is unloaded (exited), the subordinates of the parent Subscriber do not renegotiate to another parent Subscriber. The subordinates wait until that parent Subscriber is loaded again and continue to use it. The reason for this is that if the parent Subscriber was the only server between twenty Subscribers and the Distributor (which is located across the WAN), you do not want all of the Subscribers to go across the WAN to get their Distributions if the parent Subscriber is unavailable.

3.11.3 System Resources and Server Behavior

Using Policy and Distribution Services can affect the behavior of your system:

- ◆ Tiered Electronic Distribution usage can affect system behavior because of the traffic created in sending Distributions
- ◆ Some server policies are designed to control the behavior of servers, such as how a server should be brought down
- ◆ Some server policies are designed for NetWare server configuration, such as SET parameters, content of the `autoexec.ncf` file, and so on

Installing and using Tiered Electronic Distribution can affect any of the following:

- ◆ CPU utilization
- ◆ Disk space resources
- ◆ Network traffic
- ◆ Other I/O activity

To optimize your installation of Tiered Electronic Distribution, you should consider the following issues when selecting Distributor and Subscriber servers:

- ◆ Which servers are the best candidates for the heavy workload of a Distributor?

Consider CPU speed for building and sending Distributions, and sufficient disk space for storing all of the Distributor's Distributions.

The server can perform other non-Server Management network functions, be running other Server Management or non-Server Management software, or be solely dedicated to the Distributor function.

- ◆ Which servers do you want to manage using server policies?

Consider installing the Subscriber software to each server that you want to manage with policies, or where you want to distribute software packages. The policy engine is installed with the Subscriber software; also, the Subscriber software is used to extract and install software packages.

- ◆ Which servers could best handle the additional workload of being a parent Subscriber? (A parent Subscriber is a Subscriber that acts as a proxy for the Distributor to store and pass Distributions so that the Distributor does not need to send its Distributions to every Subscriber.)

Consider CPU speed for sending the Distributions, and free disk space for storing the Distributions that the parent Subscriber passes on.

- ◆ Does each of your LAN segments have servers that are capable of being a parent Subscriber?

Consider WAN traffic when deciding where to locate parent Subscribers.

- ◆ Do you have other processes using up bandwidth on some LANs and WAN links?

Consider Distribution priorities and setting sending and receiving rates to minimize the affect Distributions can have on bandwidth for WAN links.

3.11.4 Controlling I/O Rates and Concurrent Distributions

If you need to control bandwidth usage for Distribution traffic, you can set the I/O rates and the maximum number of concurrent Distributions for Distributors and/or Subscribers.

Attributes of both the Distributor and Subscriber objects provide the following controls:

- ◆ **Input rate:** For sending and receiving Distributions, you can set the maximum bytes per second. The Distributor Agent sends the Distributions, and the Policy/Package Agent receives and extracts them. This allows you to have some control over the bandwidth used by these agents. The default is the maximum that the connection can handle. However, this does not control the rate at which FTP, HTTP, and RPM Distributions are built by the Distributor.
- ◆ **Output rates based on Distribution's priority:** Sets the default output rate to minimize network traffic for Tiered Electronic Distribution objects. This determines the send rate for Distributors and parent Subscribers. The default value is the maximum that the connection can handle. Blank means that bandwidth is taken from third-party applications.

There are three output priorities where you can specify a rate:

- ◆ **High priority:** These Distributions are sent before any Medium or Low priority Distributions.
- ◆ **Medium priority:** These Distributions are sent after all High priority and before any Low priority Distributions.

- ♦ **Low priority:** These Distributions are sent after all High and Medium priority Distributions.

For more information, see [Section 3.4.5, “Prioritizing Distributions,” on page 126.](#)

- ♦ **Maximum number of concurrent Distributions:** This determines how many simultaneous Distributions the Distributor Agent can send. The default is unlimited (blank field). The Subscriber always receives as many Distributions as it is sent; however, it only concurrently passes on the number that you choose here.

If there is only one Subscriber, the Distributor sends Distributions at the selected rate. If there are two Subscribers, the Distributions are sent at one half the rate. In other words, to determine the slowest distribution rate, divide the Distributor’s output rate by the maximum number of concurrent Distributions.

Because Subscribers always receive another concurrent Distribution, the rate applies even though you cannot limit the number of incoming connections.

3.11.5 Minimizing Messaging Traffic

Tiered Electronic Distribution provides message notifications so that administrators and selected end users are kept informed. Notifications are sent in several ways:

- ♦ Information written to log files
- ♦ Notifications sent via e-mail messages
- ♦ SNMP traps used and displayed on both local and remote consoles

The following sections explain message notification usage:

- ♦ [“Message Notification Levels” on page 183](#)
- ♦ [“Managing Message Notification Level Log Files” on page 184](#)
- ♦ [“Sending Notifications Over LANs and WANs” on page 184](#)

Message Notification Levels

There are seven levels of message notification available. Each level adds its own information to the previous level.

| Messaging Level | Description |
|--|---|
| Level 0 - No messages | Messages are not sent. |
| Level 1 - Errors | Reports unusual or unexpected situations that can cause an operation to fail. They often require user intervention to correct the problem. |
| Level 2 - Successes and Level 1 messages | Reports completion of a successful operation. |
| Level 3 - Warnings and Level 2 messages | Reports unusual but not unexpected runtime conditions. These messages usually do not require user intervention, but in some situations an unusual runtime condition does. |
| Level 4 - Information and Level 3 messages | Informs the user about what has happened or is currently happening. They usually do not require any action from the user. |

| Messaging Level | Description |
|--|---|
| Level 5 - Trace and Level 4 messages | Reports detailed trace information that is used to troubleshoot unusual or unexpected situations that cause an operation to fail. This information might only be useful with the guidance of Novell Support.. |
| Level 6 - Developer trace and Level 5 messages | Used for debugging code. This information is useful only to Novell Support and Development. |

Regardless of the destination for a message, resources are directly affected by the level you choose.

For information on setting message notification levels, see:

- ◆ **Distributor object:** [Step 3 on page 107](#)
- ◆ **Subscriber object:** [Step 3 on page 151](#)

Managing Message Notification Level Log Files

The level you choose for a log file affects the rate at which the log file grows. Because log files have no maximum size, you can control the size of a log file by choosing to delete entries after *x* number of days. For information on setting message notification levels for log files, see:

- ◆ **Distributor object:** [Step 2 on page 106](#)
- ◆ **Subscriber object:** [Step 2 on page 150](#)

Sending Notifications Over LANs and WANs

The greatest impact on network traffic can come from the levels you choose for SNMP traps and for the remote console. For information on setting message levels for SNMP traps, e-mail messages, and the server's console, see:

- ◆ **Distributor object:** [Step 3 on page 107](#)
- ◆ **Subscriber object:** [Step 3 on page 151](#)

SNMP Traps

SNMP messages are sent only if there is an SNMP policy in effect for the receiving server, regardless of the level you choose for the messages. SNMP traffic is affected by both the level you choose and by the SNMP configuration in the policy on the server. There is one SNMP packet per message per destination in the SNMP Trap Target policy. IPX™ addresses are not supported for trap targets.

E-Mail Messages

E-mail messages can also affect network traffic. Like SNMP, e-mail sends only one e-mail per message per e-mail user defined. E-mail is also configured by a server policy. You must define and enable the policy on the sending server for e-mail messages to be sent.

3.11.6 Changing DNS Names or IP Addresses for Tiered Electronic Distribution Servers

Whenever there is a change to the identity of either a Distributor or Subscriber server, you must perform certain tasks so that the distribution processes for these servers can continue as before.

In the distribution process, Tiered Electronic Distribution servers identify themselves to each other by their DNS names or IP addresses. The following sections explain situations that can arise from changing these server identifiers.

If You Are Using DNS Names to Identify Your Servers

- ♦ If you change the DNS name of a Distributor server, Subscriber servers cannot recognize the Distributor as a valid source for receiving Distributions.
- ♦ If you change the DNS name of a Subscriber server, the Distributor cannot locate the Subscriber server for sending Distributions to it. This is because the Distributor obtains the Subscriber server's address from the eDirectory object.

If you change the IP address of a Distributor or Subscriber server when you are using its DNS name to identify it to Server Management, this change does not affect the distribution processes.

If You Are Using IP Addresses to Identify Your Servers

- ♦ If you change the IP address of a Distributor server, Subscriber servers cannot recognize the Distributor as a valid source for receiving Distributions.
- ♦ If you change the IP address of a Subscriber server, the Distributor cannot locate the Subscriber server for sending Distributions to it. This is because the Distributor obtains the Subscriber server's address from the eDirectory object.

Because reinstating valid certificates is involved in resolving server identity changes, see [Section 7.1.7, "Handling Invalid Certificates," on page 304](#) for instructions.

3.11.7 When a Tiered Electronic Distribution Process Fails

It is possible, for many common computer-related reasons, that a Tiered Electronic Distribution process could fail. The following are a few possibilities:

- ♦ **A Distribution could be interrupted.** If so, when it restarts it picks up where it left off.

Before distribution, the Distribution package resides at the Distributor. After distribution, the Distribution package still resides at the Distributor with a copy now at the Subscriber. It is during the distribution process that an interruption could halt copying. When the Distributor tries to re-send the Distribution (the next time the Channel schedule starts), it picks up where it left off and does not re-send the entire Distribution.

If re-sending a Distribution is interrupted, the sender retries every two minutes for 30 minutes. If it is not successful in reestablishing connection to the target server, it stops retrying. The next time the Channel's schedule starts, it picks up where it left off in sending the Distribution when it was originally interrupted.

- ♦ **An extraction could be interrupted.** If so, the extraction does not pick up where it left off.

Distributions are made across the wire from server to server, while extractions are performed on the server from Distributions already sent. Therefore, when an extraction is interrupted, it

simply fails. The Subscriber does not roll back (or undo) the failed extraction, unless the Distribution was a software package (.cpk file). It tries the extraction again the next time the Subscriber's extraction schedule starts.

Files are extracted to the volume and directory specified when the Distribution package was created. File groupings and software packages both allow you to specify to which volume and directory the package should be extracted. Therefore, when an interruption occurs during extraction, it fails in the same way as if you were copying a file in the operating system.

- ◆ **The File type offers the following:**

- Retry *X* times
- Kill the connection on files that are open
- Error handling (Fail on error; perform a routine on error)

All options deal with extraction and how to handle it.

3.12 Working Directories

Distributors and Subscribers use working directories on the servers for Distributions, patches, status files, and temporary working files. The size of a working directory is determined by the size and number of Distributions.

The working directories default to the sys: volume on NetWare servers or the C: Drive on Windows servers. Because of disk space considerations on NetWare servers, we recommend that you select a different location on the server, such as a data: volume.

The default working directory names for NetWare and Windows servers are *path\zenworks\pds\ted\dist* for the Distributor and *path\zenworks\pds\ted\sub* for the Subscriber. For Linux and Solaris servers, the paths are */var/opt/novell/zenworks/zfs/pds/ted/working/dist* and */var/opt/novell/zenworks/zfs/pds/ted/working/sub*. You can change working directory names in the properties of the Tiered Electronic Distribution object.

The following sections describe the Tiered Electronic Distribution directory structures:

- ◆ [Section 3.12.1, "NetWare Distributor Directories," on page 186](#)
- ◆ [Section 3.12.2, "NetWare Subscriber Directories," on page 188](#)
- ◆ [Section 3.12.3, "Windows Distributor Directories," on page 189](#)
- ◆ [Section 3.12.4, "Windows Subscriber Directories," on page 189](#)
- ◆ [Section 3.12.5, "Linux or Solaris Distributor Directories," on page 190](#)
- ◆ [Section 3.12.6, "Linux or Solaris Subscriber Directories," on page 190](#)

3.12.1 NetWare Distributor Directories

The following directories are used by NetWare Distributors:

volume:\installation_path\zenworks\pds\ted

Contains the Tiered Electronic Distribution software for the Distributor.

volume:\installation_path\zenworks\pds\ted\security\private

Contains the Distributor's private key.

volume:\working_directory

Contains one directory for each Distribution that belongs to the Distributor. The working directory name is user-defined in the Distributor object.

volume:\working_directory\distribution_directory

Each Distribution has its own directory that is created under the working directory. The Distribution directory's name is derived from the following syntax: *Tree_DN_of_Distribution*. For example, `TestTree_Files.Distributions.ZENworks.Novell`.

volume:\working_directory\distribution_directory\time_stamp_directory

Each Distribution directory contains multiple time-stamp directories, which are named according to the date and time the Distribution was built.

Each time a Distribution is built, the Distributor checks to see if anything has changed since the last time the Distribution was built. If so, a new time-stamp directory is created.

The number of time-stamp directories kept is determined by the Maximum Number of Revisions to Keep field in the Distribution object's properties. There are occasions when the number of time-stamp directories exceeds the maximum number specified, because the Distributor does not delete a time-stamp directory that is in use. The Distributor removes the oldest time-stamp directories first.

Sometimes a time-stamp directory name has `_temp` appended to it. When a Distributor builds a Distribution, it creates a `*_temp` directory before it determines if anything has changed. If changes are discovered, `_temp` is removed and the directory is used for the new build.

A Distributor's time-stamp directories contain the files listed in [Table 3-18](#):

Table 3-18 Files in the Distributor's Time-Stamp Directories

| Filename | Description |
|---------------------------|--|
| <code>distfile.ted</code> | The Distribution that was built. All Distributions have the same filename. They are distinguished by their time-stamp directory's name and path. |

| Filename | Description |
|--------------------|---|
| <i>digest_file</i> | <p>This file only exists if the Distributor Agent creates it (optional).</p> <p>Digests are used by Distributors and Subscribers to verify that Distributions have not been tampered with while in transit. The digest provides an MD5 checksum for the Subscriber to compare.</p> <p>Digests also detect corruption in a Distribution's package. In the case of corruption, the Subscriber renames the <i>distfile.ted</i> Distribution file to <i>distfile.corrupt</i> and the Distribution is rebuilt and sent the next time the Channel's schedule fires.</p> <p>The syntax for creating the digest filename is:</p> <pre>%AGENT%AgentDigest.ted</pre> <p>For example:</p> <pre>FTPAgentDigest.ted HTTPAgentDigest.ted FileAgentDigest.ted CPKAgentDigest.ted</pre> |

3.12.2 NetWare Subscriber Directories

The following directories are used by NetWare Subscribers:

volume:\installation_path\zenworks\pds\ted

Contains the Tiered Electronic Distribution software for the Subscriber and/or Distributor.

volume:\installation_path\zenworks\pds\ted\security

Contains certificates received from Distributors.

volume:\working_directory

Contains one directory for each Distribution that it receives from a Distributor. The working directory name is user-defined in the Subscriber object.

volume:\working_directory\distribution_directory

Each Distribution has its own directory that is created under the working directory. The Distribution directory's name is derived from the following syntax: *Tree_DN_of_Distribution*. For example, *TestTree_Files.Distributions.ZENworks.Novell*.

volume:\working_directory\distribution_directory\time_stamp_directory

Each Distribution directory contains multiple time-stamp directories, which are named according to the date and time the Distribution was built.

The number of time-stamp directories kept is determined by the Maximum Number of Revisions to Keep field in the Distribution object's properties.

After a threshold is met, the Subscriber receives the maximum revision information and deletes the oldest time-stamp directories first.

A Subscriber's time-stamp directories contain the files listed in [Table 3-19](#):

Table 3-19 Files in the Subscriber's Time-Stamp Directories

| Filename | Description |
|-----------------------------|--|
| <code>distfile.ted</code> | The Distribution that was built. All Distributions have the same filename. They are distinguished by their time-stamp directory's name and path. |
| <code>diststatus.ted</code> | After a Distribution has been successfully received, this file is created. |
| <code>digest_file</code> | This file only exists if the Distributor Agent has created it (optional). Digests are used by Distributors and Subscribers to verify that Distributions have not been tampered with while in transit. The digest provides an MD5 checksum for the Subscriber to compare. Digests also detect corruption in a Distribution's package. In the case of corruption, the Subscriber renames the <code>distfile.ted</code> Distribution file to <code>distfile.corrupt</code> and the Distribution is rebuilt and sent the next time the Channel's schedule fires. |

3.12.3 Windows Distributor Directories

The following directories are used by Windows Distributors:

installation_path\zenworks\pds\ted

Contains the Tiered Electronic Distribution software for the Distributor.

installation_path\zenworks\pds\ted\security\private

Contains the Distributor's private key.

3.12.4 Windows Subscriber Directories

The following directories are used by Windows Subscribers:

installation_path\zenworks\pds

Contains the Tiered Electronic Distribution software for the Subscriber.

installation_path\zenworks\pds\ted\security\private

Contains certificates received from Distributors.

local_drive:\working_directory\distribution_directory\time_stamp_directory

Each Distribution directory contains multiple time-stamp directories, which are named according to the date and time the Distribution was built.

3.12.5 Linux or Solaris Distributor Directories

The following directories are used by Linux or Solaris Distributors:

`/var/opt/novell/zenworks/zfs/pds/ted/working/dist`

Contains the Tiered Electronic Distribution software for the Distributor.

`/var/opt/novell/zenworks/zfs/pds/ted/security/private`

Contains the Distributor's private key.

Each Distribution directory contains multiple time-stamp directories, which are named according to the date and time the Distribution was built.

3.12.6 Linux or Solaris Subscriber Directories

The following directories are used by Linux or Solaris Subscribers:

`/var/opt/novell/zenworks/zfs/pds/ted/working/sub`

Contains the Tiered Electronic Distribution software for the Subscriber.

`/var/opt/novell/zenworks/zfs/pds/ted/security/private`

Contains certificates received from Distributors.

Each Distribution directory contains multiple time-stamp directories, which are named according to the date and time the Distribution was built.

3.13 Editing the Tednode.properties File

If you should install the Subscriber software to a server that does not have a Subscriber object in any eDirectory tree, such as a Windows server in a Microsoft domain, the `tednode.properties` file is used by such a server for its configuration information. When you have Subscriber configuration changes, you need to edit the server's `tednode.properties` file using the information in this section.

The `tednode.properties` file is located in the `\zenworks\pds\ted` directory on the server.

Table 3-20 shows the required format of the file, including comments on some of the entries. The information on the right side of an = symbol is only an example and not the required value for that line. However, the examples are intended to show the correct syntax for the values.

Table 3-20 *Tednode.Properties File Fields*

| Line Content | Comments |
|---------------------------------------|----------------------------------|
| <code>workingdir = d:\ted\tran</code> | Subscriber's working directory |
| <code>io.input = 100</code> | Receive rate in bytes per second |

| Line Content | Comments |
|--|---|
| io.output = -1 | Send rate in bytes per second |
| variable1 = vol=sys: | Define the variable "vol" with the value "sys:" |
| variable1.description = Destination Volume | A description of the variable's function |
| console.level = 6 | Message level for the server's console |
| log.level = 1 | Message level for log file |
| log.days = 1 | Number of days to save log file entries |
| log.path = d:\ted\tran\log.txt | Path for log file and log filename |
| workorder.timeout = 0 | Number of seconds to wait for reply from the Distributor before dropping connection; 0 = wait forever |
| workorder.concurrent = 0 | Concurrent Distributions |
| email.level = 0 | Message level for e-mail |
| smtp.host = email.novell.com | Location of SMTP host |
| snmp.level = 0 | Message level for SNMP traps |
| email.target1 = johndoe@novell.com | E-mail address for the messages |

For the remaining `tednode.properties` file entries, remove the # (comment) symbol from a line to enable it. This makes that line effective for the schedule type that it is listed under. However, do not remove the # symbol from the first line for a schedule type because it is only a description that indicates the schedule type. You can change the default values that are listed.

The following sample has the Daily schedule enabled because the appropriate # symbols have been removed:

Line Content

```
# Yearly schedule and associated keys (with default values specified)
#schedule.type=yearly
#schedule.month=1
#schedule.day=1
#schedule.begin.hour=8
#schedule.begin.minute=0
#schedule.end.hour=17
#schedule.end.minute=0
#schedule.random=false

# Monthly schedule and associated keys (with default values specified)
#schedule.type=monthly
#schedule.day=1
#schedule.begin.hour=8
#schedule.begin.minutes=0
#schedule.end.hour=17
#schedule.end.minute=0
#schedule.random=false
```

Line Content

Daily schedule and associated keys (with default values specified)

```
schedule.type=daily
schedule.days=Monday,Tuesday,Wednesday,Thursday,Friday,Saturday,Sunday
schedule.begin.hour=8
schedule.begin.minutes=0
schedule.end.hour=17
schedule.end.minute=0
schedule.repeat.days=0
schedule.repeat.hours=0
schedule.repeat.minutes=0
schedule.random=false
```

Immediate schedule and associated keys (with default values specified)

```
#schedule.type=immediately
#schedule.repeat.days=0
#schedule.repeat.hours=0
#schedule.repeat.minutes=0
```

Interval schedule and associated keys (with default values specified)

```
#schedule.type=interval
#schedule.repeat.hours=0
#schedule.repeat.minutes=0
```

Never schedule and associated keys (with default values specified)

```
#schedule.type=never
```

Time schedule and associated keys (with default values specified)

```
#schedule.type=time
#schedule.date.year=2001
#schedule.date.month=1
#schedule.date.day=1
#schedule.begin.hour=8
#schedule.begin.minutes=0
```

Server Policies

4

Novell® ZENworks® Server Management provides server policies for managing server configurations, processes, and behaviors.

The following sections will help you to understand, set up, and configure the policies:

- ◆ [Section 4.1, “Understanding Server Policies,” on page 193](#)
- ◆ [Section 4.2, “Creating a Policy Package,” on page 204](#)
- ◆ [Section 4.3, “Configuring Server Policies,” on page 205](#)
- ◆ [Section 4.4, “Enabling Policies,” on page 233](#)
- ◆ [Section 4.5, “Distributing Policies,” on page 233](#)
- ◆ [Section 4.6, “Associating Policies,” on page 233](#)
- ◆ [Section 4.7, “Scheduling Policies,” on page 234](#)
- ◆ [Section 4.8, “Viewing Effective Policies,” on page 234](#)
- ◆ [Section 4.9, “Changing Policy Enforcement,” on page 234](#)

4.1 Understanding Server Policies

In ZENworks 7 Server Management, most policies are enforced through the distribution of policy packages. However, a few policies used by the Distributor are enforced by being associated with Novell eDirectory™ containers.

Review the following sections to understand policies in ZENworks 7 Server Management:

- ◆ [Section 4.1.1, “Configuration and Behavioral Management through Server Policies,” on page 193](#)
- ◆ [Section 4.1.2, “Server Policies and Policy Packages,” on page 194](#)
- ◆ [Section 4.1.3, “Policy Characteristics,” on page 195](#)
- ◆ [Section 4.1.4, “Server Policies Architecture,” on page 196](#)
- ◆ [Section 4.1.5, “Enforcing Policies,” on page 199](#)
- ◆ [Section 4.1.6, “Server Policy Descriptions,” on page 200](#)

4.1.1 Configuration and Behavioral Management through Server Policies

The Server Policies component provides configuration and behavioral management of your servers. Server policies are divided into three packages for the convenience of scheduling policies and distributing the policies to their applicable servers:

- ◆ **Container Package:** Holds the Search policy that determines how Policy and Distribution Services searches eDirectory for objects associated with policies.
- ◆ **Service Location Package:** Holds policies specific to running Policy and Distribution Services.

- ♦ **Distributed Server Package:** Has a generic set of policies that can be applied to all servers, as well as policy package sets for servers on specific platforms. This package provides policies that are distributed for enforcement.

Configuration policies hold information in eDirectory that creates a similar type of configuration on a server, such as enforcing selected SET parameters. Behavioral policies hold a set of rules to be followed under certain situations, such as when a server goes down.

Through server policies you can automate the management of your servers, and through ConsoleOne® and the ZENworks Server Management role in Novell iManager you can configure policies and manage your servers from a single workstation.

4.1.2 Server Policies and Policy Packages

Server policies provide you with the ability to set, standardize, and automate configuration parameters on any given set of servers. You can control the behavior of servers in given situations, such as when downing a server.

The following sections

- ♦ [“Creating Policies” on page 194](#)
- ♦ [“Scheduling Policies” on page 194](#)
- ♦ [“All Enabled Policies Are Enforced” on page 194](#)
- ♦ [“Individual Policy Changes Are Not Tied to the Policy Package” on page 195](#)

Creating Policies

To use server policies, you must first create the appropriate Policy Package objects in ConsoleOne, configure the policies that your server needs, enable them, and distribute the package to the applicable Subscriber servers where the package’s policies are to be enforced.

Scheduling Policies

When you set up server policies, you can individually schedule them to run daily, weekly, monthly, yearly, by an event, at a specific date and time, relative to a date and time, by an interval of time, or even immediately. The default schedule for the individual policies is the default for the policy package’s schedule. Therefore, when you change the package’s default schedule, any policy in the package that doesn’t have a schedule specified then uses the package’s new schedule.

All Enabled Policies Are Enforced

You can implement (enable) any or all of the Policy and Distribution Services policies in a policy package. You can also create a Policy Package object for each different configuration set that you need. For example, you might want some of your servers to be brought down differently, so they would use different policy packages.

All policies enabled in a package are enforced on any servers where the Policy Package Distribution has been received and extracted. In other words, you cannot selectively enforce certain policies in a package. All policies in the package that are enabled are enforced on the server.

Individual Policy Changes Are Not Tied to the Policy Package

Because each policy in a policy package has its own (hidden) object in eDirectory, any changes you make to a policy that are saved when you exit the policy's dialog box (by clicking either OK or clicking Apply then Close), are not undone if you then click Cancel on the policy package's dialog box.

Therefore, clicking Cancel on the properties page for the policy package applies only to the changes you might have made for the package. For example, enabling or disabling a policy, adding or removing added policies.

Disabling a policy does not undo any configurations you made previously in the policy. The policy's configuration changes remain, but are not used because the policy is disabled.

4.1.3 Policy Characteristics

There are two different aspects of policies that determine how you use them:

- ◆ [“Plural and Cumulative Policies” on page 195](#)
- ◆ [“Configuration and Behavioral Policies” on page 195](#)

Plural and Cumulative Policies

Policy packages can contain both plural and cumulative policies. All plural policies are also cumulative, but cumulative policies are not necessarily plural. For more detail, review:

- ◆ [“Plural Policies” on page 195](#)
- ◆ [“Cumulative Policies” on page 195](#)

Plural Policies

Plural policies are those where there can be more than one per policy package per platform.

For example, in the same policy package, you can add and configure a Scheduled Down policy and name it “Scheduled Down for Time A.” Then you could add and configure another Scheduled Down policy, this time naming it “Scheduled Down for Time B.”

You can tell if a policy is plural by viewing the Policies tab and clicking Add, because all plural policies are listed in the Add dialog box.

Cumulative Policies

Cumulative policies are those that allow multiples of the same policy to be in effect when multiple policy packages are distributed to a server. For example, a Text File Changes policy distributed to Server A could be accumulated with a differently configured Text File Changes policy distributed to Server A. All of the text file changes from both policies would be effective for Server A.

Configuration and Behavioral Policies

A single configuration policy can affect the configuration of a single server or many servers. For example, you can schedule a policy to run at regular intervals to ensure that the server's configuration continues to be set correctly.

Behavioral policies hold a set of rules to be followed in certain situations. The policy engine carries out these rules, along with any of its supporting modules.

For example, the Server Down Process policy defines criteria that must be met before you bring the server down, such as:

- ◆ How soon before the server is brought down should users be notified
- ◆ Who is notified when the policy is being enforced
- ◆ Which peer server is to send SNMP alerts if the server does not come back up

IMPORTANT: For Linux and Windows servers, any downing command entered locally on those servers cannot be intercepted by the Server Down Process policy. NetWare servers use APIs that enable the policy to intercept the action. For the Server Down Process policy to work for the Linux and Windows server platforms, they must be downed using iManager where the action can be detected by the policy.

Behavioral policies are designed to make servers act more intelligently, to handle situations an administrator might not even be aware of, and to reduce complexity for administrators.

In summary, the benefits of configuration and behavioral policies include:

- ◆ Automating tasks that an administrator would normally perform
- ◆ Notifying specified users through e-mail messages that a server is going down
- ◆ Allowing a server down process to abort on certain conditions

4.1.4 Server Policies Architecture

To understand how server policies are used to manage your servers, you must understand its eDirectory objects and its agent:

- ◆ [“eDirectory Schema Extensions for Server Policies” on page 196](#)
- ◆ [“Policy/Package Agent” on page 199](#)

eDirectory Schema Extensions for Server Policies

The eDirectory schema extensions included in the Server Policies component define the class of eDirectory objects that are created in your eDirectory tree, including which information is required or optional at the time the object is created. Every object associated with the Server Policies component in an eDirectory tree has a class defined for it in the tree’s schema.

Server Management objects for the eDirectory schema are:

- Container Package
- Server Package
- Service Location Package
- Distributed Server Package
- ZENworks Database

Note the following concerning policy enforcement:

- ◆ All of the policies in the Distributed Server Package must be distributed to be enforced

- ◆ All of the policies in the Container Package, Server Package, and Service Location Package must be associated to be enforced

Existing eDirectory classes that are modified with the addition of Server Management attributes are:

Country
Group
Locality
Organization
Organizational Unit
Server

The following sections summarize the primary eDirectory objects that are added to eDirectory from the schema extensions provided with the Server Policies component:

- ◆ “Container Package Object” on page 197
- ◆ “Server Package Object” on page 197
- ◆ “Service Location Package Object” on page 197
- ◆ “Distributed Server Package” on page 198
- ◆ “ZENworks Database Object” on page 198

For basic information about the types of objects in an eDirectory tree, see the [Novell NetWare Documentation Web site \(http://www.novell.com/documentation/lg/nw5/docui/index.html\)](http://www.novell.com/documentation/lg/nw5/docui/index.html) and select *Procedures > Planning > Directory Services > eDirectory Planning*.

Container Package Object

The Container Package object is an eDirectory object that manages the Search policy object. This policy is used by the Distributor and Subscriber objects for all versions of Server Management, and must be associated to be enforced.

Server Package Object

The Server Package object is an eDirectory object that manages the following policy objects for ZENworks Server Inventory:

Rollup Policy
zeninvDictionaryUpdatePolicy
ZENworks Database

All policies in this package must be associated to be enforced.

Policy and Distribution Services does not use this package.

Service Location Package Object

The Service Location Package object is an eDirectory container object that manages the following policy objects:

SMTP Host
SNMP Trap Targets
Tiered Electronic Distribution
ZENworks Database

Service Location Package policies provide general Policy and Distribution Services configuration and location information.

All policies in this package must be associated to be enforced.

All policies are used by ZENworks 7 Server Management Distributors and Subscribers.

Distributed Server Package

The Distributed Server Package object is an eDirectory object that manages the following policy objects (ZENworks 7 Server Management only):

- Copy Files
- NetWare Set Parameters
- Prohibited File
- Scheduled Down
- Scheduled Load/Unload
- Server Down Process
- Server Scripts
- SMTP Host
- SNMP Community Strings
- SNMP Trap Targets
- Text File Changes
- ZENworks Database
- ZENworks Server Management

Distributed Server Package policies are used for configuring servers, controlling server behavior, and providing general Server Management configuration and location information.

All policies in this package must be distributed to be enforced.

ZENworks Database Object

Provides the location of the `zfslog.db` file for logging reporting information. You can install the database file on only NetWare[®] and Windows servers.

The ZENworks Database object can exist multiple times in a tree, each with its own associated database file; however, there can only be one database file installed per server.

The Server Policies component writes policy information to the Server Management database file (`zfslog.db`). Because every server in your network can be running the Policy/Package Agent, they can each write to the database, even across WAN links. If you do not need consolidated server policies reports on all servers, you can install a database to each WAN segment.

If you require consolidated server policies reports, you can have just one `zfslog.db` file where all servers running the Policy/Package Agent can log information. The amount of data a Policy/Package Agent writes to the database might not create excessive WAN traffic, depending on the number of servers and speeds of the WAN links.

Because you can install the Server Management database to multiple servers, to minimize WAN traffic you should coordinate the placement of Policy Package and ZENworks Database objects in containers on the WAN segments.

Policy/Package Agent

Policy and Distribution Services allows you to manage your network servers using the Policy/Package Agent. This agent is installed on each server where you select the Subscriber/Policies installation option.

The Policy/Package Agent does the following:

- ◆ Extracts (installs) a software package's contents.
- ◆ Extracts the policy information from a Policy Package Distribution.
- ◆ Enforces the enabled policies from the extracted policy information based on their enforcement schedules.

There are a number of server policies that provide configuration and behavioral management of your servers. The Policy/Package Agent must be running on each server you want to manage with policies or have software packages to extract and install.

You should install the Policy/Package Agent to every server in your network. Exceptions might be servers where you do not need to distribute software packages, or servers that you do not want to manage using policies.

4.1.5 Enforcing Policies

Most ZENworks 7 Server Management policies are enforced by creating the policy package, enabling and configuring the policy, scheduling the package, distributing the package, and extracting the policies on servers.

Some ZENworks 7 Server Management policies are enforced by creating the policy package, enabling and configuring the policy, scheduling the package, and associating the package with the containers where the Distributor or Subscriber objects reside.

For more information, review the following:

- ◆ [“Scheduling Policies” on page 199](#)
- ◆ [“Distributing Policies” on page 200](#)
- ◆ [“Associating Policies” on page 200](#)

Scheduling Policies

Some server policies must be scheduled before they can be enforced.

The following schedules are available:

- ◆ Activate by the Default Package Schedule (which you can set to any of the schedules)
- ◆ Activate on a specified event (such as running at system startup or shutdown)
- ◆ Activate once relative to a period of time
- ◆ Activate at a specified date and time
- ◆ Activate once per year at a specified time
- ◆ Activate once each month at a specified time
- ◆ Activate on one or more days of the week at specified times
- ◆ Activate on one or more days of the week, repeating at a specified interval of time

- ◆ Continuously repeat at a specified interval of time
- ◆ Run immediately
- ◆ Run immediately, repeating at a specified interval of time

IMPORTANT: If you enable a policy, but do not schedule it, it activates according to the schedule currently specified in the Default Package Schedule.

The Default Package Schedule provides a default for unscheduled policies in the policy package. The default schedule is the Run At System Startup event.

Distributing Policies

After you have enabled and configured a policy contained in the Distributed Server Package, you must distribute its policy package to the Subscriber servers where the enabled policies are placed into effect. In other words, configuring and enabling a policy only sets up the policy. It is enforced through its distribution to and extraction on the applicable servers that are running Policy and Distribution Services.

Associating Policies

After you have enabled and configured a policy contained in the Service Location Package, you must associate its policy package with the containers where Distributor or Subscriber objects reside so that the enabled policies are placed into effect. This association can be directly with a container where the Distributor or Subscriber objects reside, or with a container higher in the tree from where the container holding these objects reside.

Because configuring and enabling a policy only sets up the policy, it is enforced through its association with the applicable servers that are running Policy and Distribution Services.

4.1.6 Server Policy Descriptions

The tables in the following sections list the server policies by policy package. The second column indicates whether a policy is a configuration or behavioral policy, and whether it is cumulative, plural, or both.

- ◆ [“Container Package” on page 201](#)
- ◆ [“Service Location Package” on page 201](#)
- ◆ [“Server Package” on page 202](#)
- ◆ [“Distributed Server Package” on page 203](#)

Container Package

Table 4-1 *Container Package Policy*

| Policy Name | Policy Type Keys | Policy Function |
|-------------|------------------|---|
| Search | Behavioral | <p>If you don't set a Search policy, the default is to search from the parent container to the root every hour. This can create unnecessary search traffic. Therefore, we recommend that you make effective use of the Search policy.</p> <p>This Search policy can only be administered in ConsoleOne. A Search policy created in NetWare Administrator for ZENworks is not recognized in Server Management.</p> |

Because most policies in Server Management are distributed rather than associated for enforcement and a Distributor does not receive Distributions, the Search policy is used in Server Management to enable the Distributor Agent to locate and use policies in the Service Location Package. For example, the Distributor Agent can use the package's ZENworks Database policy to write reporting information to the ZENworks Server Management Database file.

Also, Distributors read the Service Location Package policies for their Subscribers. That means Subscribers receive their Service Location Package policies through associations, as well.

Service Location Package

Table 4-2 *Service Location Package Policies*

| Policy Name | Policy Type Keys | Policy Function |
|-------------------|------------------|---|
| SMTP Host | Configuration | <p>Sets the TCP/IP address of the relay host that processes outbound Internet e-mail. This policy must be enabled if you select the E-Mail option for notifying or logging messages in any of the other policies.</p> |
| SNMP Trap Targets | Configuration | <p>Sets SNMP trap targets for associated eDirectory objects.</p> <p>In ZENworks 7 Server Management, you can schedule this policy for when you want it to be refreshed.</p> <p>IPX™ addresses are not supported for SNMP trap targets. You can only use IP addresses and DNS names.</p> |

| Policy Name | Policy Type Keys | Policy Function |
|--------------------------------|------------------|--|
| Tiered Electronic Distribution | Configuration | <p>Sets defaults for the Distributor and Subscriber objects, including:</p> <ul style="list-style-type: none"> I/O rates Maximum concurrent Distributions Connection time-out in minutes Working directory Parent Subscriber Messaging levels for a server's console, SNMP traps, log files, and e-mail notification Extraction Schedule Refresh Schedule Variables <p>Any defaults set here override unchanged defaults in a Tiered Electronic Distribution object. However, if a Tiered Electronic Distribution object's properties are modified, those modifications have precedence over any defaults set in the Tiered Electronic Distribution policy.</p> |
| ZENworks Database | Configuration | <p>Sets the DN for locating the ZENworks Database object and the database file. The database is used for logging successes and failures that are used in creating reports.</p> <p>This policy can be created to override the database settings that might have been established during installation of Policy and Distribution Services.</p> <p>The Policy/Package Agent and the Distributor Agent both write to <code>zfslog.db</code>. For information on having these agents write to different database files, see Section 10.1.7, "Coexisting Databases," on page 353.</p> |

Server Package

The Server Package exists in ZENworks 7 Server Management only for use by Server Inventory. The ZENworks Database policy contained in this package is automatically created by the installation program when Server Inventory is installed to enable automatic location of the database for logging inventory data.

Policy and Distribution Services does not use this package.

Although other policies exist in this package, [Table 4-3](#) only lists the ZENworks Database policy.

Table 4-3 *Server Package Policy*

| Policy Name | Policy Type Keys | Policy Function |
|-------------------|------------------|--|
| ZENworks Database | Configuration | Sets the DN for locating the ZENworks Database object. This policy must be in effect for Server Inventory to locate a database for logging inventory data. |

Distributed Server Package

This package contains the policies that must be distributed to Server Management servers to be enforced on them.

Table 4-4 *Distributed Server Package Policies*

| Policy Name | Policy Type Keys | Policy Function |
|------------------------|---|--|
| Copy Files | Plural Cumulative Configuration | Enables copying of files on a server from one location to another by using policy configurations. |
| NetWare Set Parameters | Plural Cumulative Configuration | Specifies and optimizes selected Set Parameters for a server or group of servers. For the NetWare platform only. |
| Prohibited File | Plural Cumulative Configuration | Monitors and enforces the deletion or moving of unauthorized files from a specified volume/drive or directory/folder. |
| Scheduled Down | Plural Cumulative Configuration Behavioral | Schedules when a server should go down, and whether it should be automatically brought back up. The policy includes which command to use in bringing it down (RESET, RESTART, or DOWN). |
| Scheduled Load/Unload | Plural Cumulative Configuration | For automating the loading and unloading order of NLM™ and Java Class processes for the selected servers, and for starting and stopping Windows services. NLM files that require user input to unload cannot be automated. |
| Server Down Process | Behavioral | For controlling which processes to follow and which conditions to meet before downing a server. |
| Server Scripts | Plural Cumulative Configuration | For automating script usage on your servers. |
| SMTP Host | Configuration | Sets the TCP/IP address of the relay host that processes outbound Internet e-mail. This policy must be enabled if you select the E-Mail option for notifying or logging messages in any of the other policies. |
| SNMP Community Strings | Configuration | Allows you to receive and respond to SNMP requests. |
| SNMP Trap Targets | Configuration | Sets SNMP trap targets for associated eDirectory objects. You can schedule this policy for when you want it to be refreshed. IPX addresses are not supported for SNMP trap targets. You can only use IP addresses and DNS names. |
| Text File Changes | Plural Cumulative Configuration | For automating changes to text files. |

| Policy Name | Policy Type Keys | Policy Function |
|----------------------------|------------------|---|
| ZENworks Database | Configuration | <p>Sets the DN for locating the ZENworks Database object and the database file. The database is used for logging successes and failures that are used in creating reports.</p> <p>This policy can be created to override the database settings that might have been established during installation of Policy and Distribution Services.</p> <p>The Policy/Package Agent and the Distributor Agent both write to <code>zfslog.db</code>. For information on having these agents write to different database files, see Section 10.1.7, “Coexisting Databases,” on page 353.</p> |
| ZENworks Server Management | Configuration | <p>Basic configuration parameters for Policy and Distribution Services, such as status logging, defining the server console prompt for the Policy/Package Agent, setting its working path, and setting a database purging limit.</p> <p>You can enable this policy on each server where you want to enforce server policies. However, if you do not enable the policy, Policy and Distribution Services works from pre-programmed defaults.</p> |

4.2 Creating a Policy Package

Policy and Distribution Services groups its server policies into three Policy Package objects:

- ◆ Container Package
- ◆ Service Location Package
- ◆ Distributed Server Package (ZENworks 7 Server Management only)

You can place policy packages anywhere in the tree. For ease of management, we recommend that you create an OU container for grouping the policy packages. For example, Policies.

However, if you install ZENworks Desktop Management to your tree, you could keep the Server Management and Desktop Management policies in separate containers, such as Server_Policies and Desktop_Policies.

IMPORTANT: If you have partitions that are accessed across a WAN, make sure that the Policy Package objects are in the same partition as the Server object to ensure that the Policy/Package Agent loads. Also make sure that the Search policy does not require searching outside the partition where the Server object exists.

To determine which Policy Package objects to create, first determine which policies you need.

To create Policy Package objects, review the instructions in the following sections:

- ◆ [Section 4.2.1, “Creating a Policies Container,”](#) on page 205
- ◆ [Section 4.2.2, “Creating a Policy Package Object,”](#) on page 205

4.2.1 Creating a Policies Container

To create the OU container object for holding your Policy Package objects:

- 1 In ConsoleOne, right-click the container where you want the policies container located.

IMPORTANT: Where you create the OU, and how many characters you use to name it, directly affects the number of characters that you have available for naming the plural policies. eDirectory has a 64-character limit for the full name and path in the tree for a policy.

Because you can have many different versions of one plural policy in a single policy package, you want to be able to name them descriptively. Therefore, place the OU as high in the tree as is logical, and give it a short name to provide as many characters as possible for naming the policies.

- 2 Click *New > Object*, then select *Organizational Unit*.
- 3 Provide a name for the OU, then click *OK*.

4.2.2 Creating a Policy Package Object

To create a Policy Package object:

- 1 In ConsoleOne, right-click the container you created for the Policy Package objects, click *New*, then select *Policy Package*.

The Policy Package Wizard opens.

- 2 Under *Policy Packages*, select a policy package, then click *Next*.

Available packages include: Container, Server, Service Location, and Distributed Server.

- 3 Provide a name for the package, then click *Next*.

Because you can have multiples of the same package type, use a unique, informative name for each package.

IMPORTANT: Because of the eDirectory 64-character path/name limit, and the package name you provide here is part of the path for plural policies that you can create later, provide a brief, but unique, Policy Package object name so that you can have as many characters as possible to be available for giving descriptive plural policy names.

- 4 Repeat **Step 2** and **Step 3** for each package to be created.

Select the *Create Another Policy Package* check box to save repeating **Step 1**.

4.3 Configuring Server Policies

You can configure server policies for containers, servers, and service locations. The policies allow you to automate use of NetWare functionality. See your NetWare documentation for specific information.

To configure server policies, review the instructions in the following sections:

- ♦ [Section 4.3.1, “Compiling Zentrapp.mib,” on page 206](#)
- ♦ [Section 4.3.2, “Configuring the Container Package Policy,” on page 206](#)

- ◆ [Section 4.3.3, “Configuring Service Location Package Policies,” on page 207](#)
- ◆ [Section 4.3.4, “Configuring Distributed Server Package Policies,” on page 214](#)
- ◆ [Section 4.3.5, “Creating Custom Log Files Using Policies,” on page 232](#)

For information on scheduling server policies, see [Section 4.7, “Scheduling Policies,” on page 234](#).

4.3.1 Compiling Zentrap.mib

The SNMP Community Strings and SNMP Trap Targets policies utilize SNMP. `Zentrap.mib` is located on the *Program CD* under `\zfs\tedpol\sfiles\mibs`.

To receive SNMP traps on your SNMP management console, you must copy the `zentrap.mib` file from the *ZENworks 7 Server Management with Support Pack 1 Program CD* to the location that your management console uses to manage MIBs, then compile it. Your SNMP management console can then receive and interpret SNMP traps from Server Management.

4.3.2 Configuring the Container Package Policy

The Search policy is used by the Distributor for information on how to read the eDirectory tree when the Distributor has been refreshed.

IMPORTANT: If you do not use the Search policy, Server Management searches up to [Root] and reads the objects every hour. Be sure to configure and enable the Search policy to limit unnecessary search traffic.

To configure the Search policy:

- 1 In ConsoleOne, right-click the *Container Package*, click *Properties*.
- 2 Select the *Policies* tab, select the check box for *Search Policy*, click *Properties*, then select the *Search Level* tab.

If the box under the *Enabled* column is not selected for the Search policy, select it before clicking *Properties*. A policy must be enabled to activate the *Properties* button.

- 3 To determine the upper limits of the search policy, select one of the following:

| Search Location | Description |
|--------------------|---|
| Object Container | Search to the parent container of the Server object |
| Partition | Search to the Partition Root |
| Selected Container | Search to the selected container |
| [Root] | Search to the root of the tree |

If you chose *Selected Container*, browse to select the container.

To determine searching limits in either direction of the item selected, enter a number. For example:

| # | Description |
|----|--|
| 0 | Limits the search to the current level (as set in the Search For Policies Up To field). |
| 1 | Limits the search to one level above the current level (as set in the Search For Policies Up To field). For example, if you specify the server's parent container in the Selected Container field, +1 would limit the search to one level above the parent container. |
| -1 | Limits the search to one level below the chosen search level (as set in the Search For Policies Up To field). For example, if you select [Root] in the Search For Policies Up To field, -1 would allow searching up to one level below [Root]. |

- 4 To determine the search order, select the *Search Order* tab.

| Type | Description |
|-----------|----------------------|
| Object | Server |
| Group | Server Group |
| Container | Container of Servers |

Use the arrow keys to change the order. You can also click *Add* or *Remove* to change which object types are used.

- 5 (Optional) Because policies are refreshed when they are received at the Subscriber, specify a refresh frequency.

The default is once every hour.

If you leave both time increments at zero (days and hours), policies are not refreshed from eDirectory, even if you have *Policy Manager Will Refresh Policies From eDirectory* selected.

Changes made to enabled policies are not enforced until they are refreshed at the given refresh interval. However, you can manually refresh all policies using the POLICY REFRESH command at the server console. The refresh rate is listed in seconds at the server console (1 hour = 3600 seconds).

- 6 Click *OK* to close the policy.

If you click *Cancel*, none of the Search policy changes made on any of the tabs are saved.

- 7 To associate the policy package so that the Search policy is enforced on the Distributor, select the *Associations* tab, then click *Add*.

- 8 Browse to select the container where the Distributor object resides (or any container above it), then click *OK*.

If you click *Cancel*, the association you made is not saved.

4.3.3 Configuring Service Location Package Policies

Because the Distributor does not receive Distributions, policies for a Distributor must be associated with the container where its object resides. The Service Location Package contains policies used by the Distributor.

To configure Service Location Package policies, review the following sections:

- ◆ “SMTP Host” on page 208
- ◆ “SNMP Trap Targets” on page 208
- ◆ “Tiered Electronic Distribution” on page 209
- ◆ “ZENworks Database” on page 214

SMTP Host

Sets the TCP/IP address of the SMTP relay host that processes outbound Internet e-mail. This policy must be enabled if you select the E-Mail option for notifying or logging messages for the Distributor.

To configure the SMTP Host policy:

- 1 In ConsoleOne, right-click the Service Location Package, then click *Properties*.
- 2 Select the SMTP Host policy’s check box, then click *Properties*.
If the box under the *Enabled* column is not selected for the SMTP Host policy, select it before clicking *Properties*. A policy must be enabled to activate the *Properties* button.
- 3 Provide the TCP/IP address or DNS name of the relay host server, then click *OK*.
- 4 To associate the policy package so that the SMTP Host policy is enforced on the Distributor, select the *Associations* tab, then click *Add*.
- 5 Browse to select the container where the Distributor object resides (or any container above it), then click *OK*.
If you click *Cancel*, the association you made is not saved.

SNMP Trap Targets

Use this property page to establish the targets (or locations) where you want SNMP traps sent from the Distributor. Each target must be a valid TCP/IP address or DNS name.

Make sure that you have compiled `zentrap.mib` (see [Section 4.3.1, “Compiling Zentrap.mib,”](#) on page 206).

To configure the SNMP Trap Targets policy:

- 1 In ConsoleOne, right-click the *Service Location Package*, then click *Properties*.
- 2 Select the SNMP Trap Targets policy, then click *Properties*.
If the box under the *Enabled* column is not selected for the SNMP Trap Targets policy, select it before clicking *Properties*. A policy must be enabled to activate the *Properties* button.
- 3 To add items to the *SNMP Trap Targets* list on the *SNMP Trap Policy* tab, click *Add*.
- 4 On the SNMP Target dialog box, provide valid a TCP/IP address or DNS name, then click *OK*.
- 5 Repeat [Step 3](#) and [Step 4](#) for each trap target to be added.
- 6 To schedule the policy, select the *Schedule* tab, select a type in the *Schedule Type* field, then configure the schedule:
[Section B.1, “Daily,”](#) on page 400
[Section B.2, “Event,”](#) on page 400

- Section B.3, “Interval,” on page 400
- Section B.4, “Monthly,” on page 401
- Section B.5, “Never,” on page 401
- Section B.6, “Package Schedule,” on page 401
- Section B.7, “Relative,” on page 402
- Section B.8, “Run Immediately,” on page 402
- Section B.9, “Time,” on page 402
- Section B.10, “Weekly,” on page 403
- Section B.11, “Yearly,” on page 403

7 Click *OK* when finished.

8 To associate the policy package so that the SNMP Trap Targets policy is enforced on the Distributor, select the *Associations* tab, then click *Add*.

9 Browse to select the container where the Distributor object resides (or any container above it), then click *OK*.

If you click *Cancel*, the association you made is not saved.

Tiered Electronic Distribution

This policy allows you to set default values for the attributes of Distributors and Subscribers.

- ◆ “How the Policy Works” on page 209
- ◆ “Cumulative Policies” on page 209
- ◆ “Replacing, Adding, or Losing Property Values” on page 210
- ◆ “Multiple Policies for Platform Configurations” on page 210
- ◆ “Configuring the Tiered Electronic Distribution Policy” on page 210

How the Policy Works

The default values set in the Tiered Electronic Distribution policy become effective when you associate the Service Location Package that contains this policy to a container above where the Distributor and Subscriber objects reside, or to the container where Subscriber objects reside.

The values in the attributes of the Tiered Electronic Distribution policy automatically replace the similar values for the Distributor and Subscriber objects, but only if the default values of those attributes have never been changed in the object’s properties.

After you have changed the values of the attributes in the Distributor or Subscriber objects and you want to use the values in the Tiered Electronic Distribution policy, then you must edit the Distributor or Subscriber object’s properties and select the Use Policy check box at the top of each tab in the object’s properties that contains the check box. Then the Tiered Electronic Distribution policy values will appear in the Distributor or Subscriber object’s attributes.

Cumulative Policies

Tiered Electronic Distribution policies are not cumulative, meaning:

- ◆ **One at a time:** You cannot have more than one Service Location Package (containing the Tiered Electronic Distribution policy) associated to the same container.

- ◆ **Closest wins:** If the Subscriber's container is associated with a Tiered Electronic Distribution policy (in the Service Location Package) and a parent container also has a Tiered Electronic Distribution policy (in the Service Location Package) associated with it, the Tiered Electronic Distribution policy of the closest container (the Subscriber's own container) prevails.

Replacing, Adding, or Losing Property Values

The following information applies only where the Tiered Electronic Distribution policy is in effect:

- ◆ You can add the Variables defined in the Tiered Electronic Distribution policy to the Distributor or Subscriber's list of variables. They do not replace the variables already defined in the Distributor or Subscriber object.
- ◆ For all other policy fields that coincide with values in a Distributor's or Subscriber's properties, the Tiered Electronic Distribution policy replaces, not supplements, them, including the possibility of replacing property values with empty fields. Therefore, if you create a Tiered Electronic Distribution policy, make sure you fill in all of the fields on every tab in the policy that you want to be applied to the affected Distributors or Subscribers.

For example, if your Subscriber has a working directory entered in its object's properties, and you do not provide a working directory in the Tiered Electronic Distribution policy, then later apply the policy by selecting the Use Policy check box on the Subscriber's properties, the Subscriber will no longer have a working directory available to it.

Multiple Policies for Platform Configurations

You can have multiple instances of the Tiered Electronic Distribution policy for your Subscriber objects for the purpose of defining different policy settings for different server platforms. To do this, you must have created the Subscriber objects in different containers representing their respective operating systems.

Subscriber attributes that could require operating system-specific values are:

- working directories
- messaging settings
- variables definitions

Configuring the Tiered Electronic Distribution Policy

1 In ConsoleOne, right-click the *Service Location Package*, then click *Properties*.

2 Select the Tiered Electronic Distribution policy, then click *Properties*.

If the box under the *Enabled* column is not selected for the Tiered Electronic Distribution policy, select it before clicking *Properties*. A policy must be enabled to activate the *Properties* button.

3 Click *General > Settings* and fill in the fields:

Input rate: Sets the default input rate to minimize network traffic for Tiered Electronic Distribution objects. This determines the receive rate for Subscribers and Distributors. The default value is the maximum that the connection can handle. You can use this rate to control the use of narrow bandwidth links.

Output rate: Sets the default output rate to minimize network traffic for Tiered Electronic Distribution objects. This determines the send rate for Distributors and parent Subscribers. The default value is the maximum that the connection can handle. Blank means that bandwidth is taken from third-party applications.

There are three output priorities where you can specify a rate:

- ♦ **High priority:** These Distributions are sent before any Medium or Low priority Distributions.
- ♦ **Medium priority:** These Distributions are sent after all High priority and before any Low priority Distributions.
- ♦ **Low priority:** These Distributions are sent after all High and Medium priority Distributions.

For more information, see [Section 3.4.5, “Prioritizing Distributions,” on page 126](#).

Maximum concurrent Distributions to build: Specifies the maximum number of distribution threads that can be running concurrently for building Distributions. The default value is 5. Valid values are from 1 to 10.

This number can help in load-balancing a Distributor’s building activity.

Maximum concurrent Distributions to send: Specifies the maximum number of distribution threads that can be running concurrently for sending Distributions. The default value is unlimited (a blank field).

This number can help in load-balancing a Distributor’s sending activity and spread network traffic over an entire scheduling window.

Connection time-out: Specifies a default number of seconds before the Distributor times out when connecting to another node, or specifies the number of seconds a Subscriber waits for a response from a Distributor (receiving) or a Subscriber (sending) before ending the connection.

After the time has transpired, a Distributor ends the connection and does not retry until the Channel’s Send schedule starts again. If a connection is ended during sending or receiving, a Subscriber does not start again until the next time the Channel’s Send schedule starts.

The default value is 300 seconds (five minutes). The available range in seconds is 1 to 60,000. You should select a reasonable time to wait for a response from one node to another.

IMPORTANT: This interval must be increased on slow or busy links where longer delays are frequent.

Working directory: Provide a default Tiered Electronic Distribution directory to store Distributions, persistent status, and temporary files on a server. The directory needs to be located where there is enough free space to handle processing of Distributions.

The Working Directory field allows the use of variables to specify the volume/drive and directory names. However, variables only work with Subscribers.

IMPORTANT: Distributors are not able to resolve variables and use exactly what is specified in the Working Directory field. For example, if the value was %VOL%ted1\working, the Distributor would create a working directory on the sys: volume named `sys:\%VOL%\ted\working`, because it could not resolve %VOL%.

For more information, see [Section 3.12, “Working Directories,” on page 186](#).

Parent Subscriber: Subscribers should generally not receive their Distributions directly from a Distributor. You can browse for a Subscriber to be the default parent Subscriber for your whole network, which passes on Distributions when a Subscriber object might not have a parent Subscriber defined in its properties.

Disk space desired to be left free: Use this as the default value to ensure there is enough free disk space for receiving Distributions where you might not have this value defined in a

Subscriber object's properties. A Subscriber does not attempt to receive a Distribution if the disk space value set here is insufficient.

4 Click *General > Messaging* and fill in the fields:

Server console: Procedure to follow when displaying messages at the server console. The default is Level 4 (Information & Level 3 Messages).

SNMP trap: Procedure to follow when sending SNMP traps. The default is Level 0 (No Messages).

Log file: Procedure to follow when recording information to a log file. The default is Level 5 (Trace Information & Level 4 Messages).

Filename: By default, this field is blank. Whatever log filename you select, it replaces `ted.log` for the servers where this policy is enforced.

To create a log file, specify the log file's filename using the following format:

`installation_path\directory_path\filename.filename_extension`

The `installation_path` is not required for ZENworks to locate the log file, but it is easier for you to locate the file if the path is included.

IMPORTANT: Because the log file can become quite large, for NetWare servers we recommend that you do not use the `sys:` volume.

Use filename extensions such as `.log` or `.txt`.

Delete log entries older than __ days: Controls disk space usage. For log files, it is important to set the message levels at minimal detail and to purge entries older than six days (the default).

E-mail: Procedure to follow when sending e-mail messages. None or Errors Only are recommended to minimize unnecessary e-mail traffic. The default is Level 0 (No Messages).

Users: Add users, groups, or e-mail addresses.

Address attribute: Displays the attribute of the associated user or group. You can change the attribute from the drop-down list, which displays over three dozen options.

Following are some of these options:

| | | |
|------------------------|-------------------------------|-------------------|
| CN | Given Name | Postal Code |
| Description | Initials | Postal Office Box |
| EMail Address | Internet EMail Address | Surname |
| Full Name | Mailbox ID | Telephone Number |
| Employee ID | NSCP:mailHost | Title |
| Entrust:User | OU | uniqueID |
| Generational Qualifier | Physical Delivery Office Name | |

5 To assign default values to variables used by the Subscriber, select the *Variables* tab, click *Add*, then fill in the fields:

Variable: Name of the variable. It should indicate how the variable is used. For example, `WORKINGVOL`.

The variable name can be derived from predefined and user-defined variables.

Value: The value that the Subscriber uses when this variable is specified. For example, `data:`.

A value can be another variable name. You can nest variables using this method.

To ensure that extraction takes place, provide an absolute path to the Subscriber. For example, if the path is only the data: volume, make sure the colon (:) is included, because it is a necessary part of the full path.

Description: Describes how the variable is used. For example:

```
Volume for the working directory.
```

If a variable defined here does not exist in a Subscriber's variables list, it is automatically added. However, if the variable does exist in the Subscriber's variables list, the definition in the Subscriber prevails.

- 6 To assign a default refresh schedule for all Distributors, select the *Schedule* tab, click *Distributor Refresh Schedule*, select a schedule in the *Schedule Type* field, then configure the schedule:

[Section B.1, "Daily," on page 400](#)

[Section B.3, "Interval," on page 400](#)

[Section B.4, "Monthly," on page 401](#)

[Section B.5, "Never," on page 401](#)

[Section B.9, "Time," on page 402](#)

[Section B.11, "Yearly," on page 403](#)

For information on the refresh schedule, see ["Scheduling" on page 96](#).

IMPORTANT: We recommend the Distributor's Refresh schedule be daily, unless changes to Distributions warrant a more frequent refresh. However, do not refresh the Distributor more often than every five minutes. The following can need up to five minutes to complete their processes: Distribution building, eDirectory replication, and tree walking (when no Search policy is defined).

- 7 To assign a default extraction schedule for all Subscribers, select the *Schedule* tab, click *Subscriber Extract Schedule*, select a schedule in the *Schedule Type* field, then configure the schedule:

[Section B.1, "Daily," on page 400](#)

[Section B.3, "Interval," on page 400](#)

[Section B.4, "Monthly," on page 401](#)

[Section B.5, "Never," on page 401](#)

[Section B.8, "Run Immediately," on page 402](#)

[Section B.9, "Time," on page 402](#)

[Section B.11, "Yearly," on page 403](#)

For information on the extraction schedule, see ["Scheduling" on page 149](#).

- 8 Click *OK* to close the policy.
- 9 To associate the policy package so that the Tiered Electronic Distribution policy is enforced on the Distributor, select the *Associations* tab, then click *Add*.
- 10 Browse to select the container where the Distributor object resides (or any container above it), then click *OK*.

If you click *Cancel*, the association you made is not saved.

- 11 To associate the policy package so that the Tiered Electronic Distribution policy is enforced on a Subscriber, select the *Associations* tab, then click *Add*.

- 12 Browse to select the container where Subscriber objects reside (or any container above it), then click *OK*.

This should be the Subscribers where you want the Tiered Electronic Distribution policy's default information to be available.

If you are creating this policy for a particular operating system, make sure you select the correct platform-specific container, and the policy applies only to the Subscribers under that container.

If you click *Cancel*, the association you made is not saved.

- 13 Repeat **Step 12** for each container where Subscribers exist that you want to use this policy.

ZENworks Database

Sets the DN for locating a ZENworks Database object. If you did not establish this information when installing Policy and Distribution Services, you can create this policy to enable Server Management to locate a database file for logging successes and failures that are used in creating reports. You can also create this policy to override the information established during installation.

Use this property page to select the database object to be associated with the current ZENworks Database policy. The policy is not in effect until you have distributed the policy to the Subscribers, or associated the policy with the Distributor.

The Server Management database is used to store reporting information for Distributions and Server Policies.

To configure the ZENworks Database policy:

- 1 In ConsoleOne, right-click the *Service Location Package*, then click *Properties*.
- 2 Select the ZENworks Database policy, then click *Properties*.
If the box under the *Enabled* column is not selected for the ZENworks Database policy, select it before clicking *Properties*. A policy must be enabled to activate the *Properties* button.
- 3 Select the *Policy/Distribution Management* tab.
- 4 In the *Database DN* field, browse for the ZENworks Database object that represents the database for this policy, then click *OK*.
- 5 To associate the policy package so that the ZENworks Database policy is enforced on the Distributor, select the *Associations* tab, then click *Add*.
- 6 Browse to select the container where the Distributor object resides (or any container above it), then click *OK*.
If you click *Cancel*, the association you made is not saved.

4.3.4 Configuring Distributed Server Package Policies

You can configure Distributed Server Package policies to automate control of various server behaviors and processes and to automate control of SMTP Host TCP/IP addresses, SNMP Trap Targets, and the ZENworks Database object's DN.

There are several Policies tab options for server policies, one for each supported operating system. The policies that are available on the General tab apply to servers on all platforms. The policies available on the specific platform tabs apply only to the servers for those platforms.

Platform-specific policies, such as those on the NetWare tab, always override similar policies on the General tab for a particular policy package.

All policies are contained in the NetWare policies. Therefore, only the NetWare policies are documented here. The information applies equally to each platform.

To configure Distributed Server Package policies, review the following sections:

- ♦ “Copy Files” on page 215
- ♦ “NetWare SET Parameters” on page 216
- ♦ “Prohibited File” on page 217
- ♦ “Scheduled Down” on page 220
- ♦ “Scheduled Load/Unload” on page 220
- ♦ “Server Down Process” on page 221
- ♦ “Server Scripts” on page 223
- ♦ “SMTP Host” on page 224
- ♦ “SNMP Community Strings” on page 224
- ♦ “SNMP Trap Targets” on page 228
- ♦ “Text File Changes” on page 229
- ♦ “ZENworks Database” on page 230
- ♦ “ZENworks Server Management” on page 231

Copy Files

The Copy Files policy enables copying of files on a server from one location to another by using policy configurations. You can either copy or move the files.

To configure the Copy Files policy:

- 1 In ConsoleOne, select the Distributed Server Package’s container, right-click the *Distributed Server Package*, then click *Properties*.
- 2 Select the *Policies* tab, then select the platform from:
 - General
 - Windows
 - NetWare
 - Linux
 - Solaris
- 3 Click *Add*, click *Copy Files*, provide a policy name, then click *OK*.
- 4 Click *Properties*.

The *Copy Files* tab displays.
- 5 Click *Add*.

Local File Copy #1 defaults. You can edit that name.
- 6 Fill in the fields:

Source path: Provide the full path where the files to be copied are located.

You can use wildcards in the path:

- * = any number of characters
- ? = any single character in that position
- ??? = any characters in those positions

Target path: Provide the full path where the copied files are to be placed.

You can use wildcards in this path. This path does not need to mirror the source path. However, you could mirror an existing target path.

Include subdirectories: Includes all subdirectories and their files beginning from the directory at the end of the path; otherwise, only the files in the directory at the end of the path are copied.

Maintain attributes: Maintains the file attributes in the target's file system that exist in the source's file system.

Overwrite destination files: Overwrites files of the same name in the destination directories, regardless of differences in file dates. If you do not select this option, files of the same name is not replaced.

Maintain trustees: Maintains the file's trustee attributes.

When a file is locked: Select one or both:

- ♦ **Retry __ times:** Retries overwriting a locked file the number of times you select before failing to replace the file. Leave this check box deselected to not replace locked files on the target file system.
- ♦ **Kill connection of open files:** (NetWare only) Attempts to kill the connection of locked files so they can be overwritten. This applies only to files being extracted, not to files being accessed to build the Distribution. If a file belonging to a Distribution is locked when the Distribution is being built, the build fails. Server and NLM connections cannot be killed.

Error processing: Fail On Error is selected by default. This stops the file copying process when an error is encountered in copying. To continue file copying when an error is encountered, select Continue On Error.

Operation: Sets whether to copy or move the files identified in the Source Path.

7 Select the *Schedule* tab, then schedule the policy (see [Section 4.7, "Scheduling Policies,"](#) on [page 234](#)).

8 Click *OK* to close the policy.

NetWare SET Parameters

You can automate the use of SET parameters by your servers.

To configure NetWare SET parameters:

- 1** In ConsoleOne, select the Distributed Server Package's container, right-click the *Distributed Server Package*, then click *Properties*.
- 2** Click *Policies > NetWare* (or *General*).
- 3** Click *Add*, then select *NetWare Set Parameters*.
- 4** Provide a name for this SET parameters policy, then click *OK*.

Because the policies selected from this dialog box are plural, you can have multiple SET parameter policies listed on the Policies tab. Therefore, provide a unique name for this policy.

When you click OK after naming the SET parameters policy, it is selected on the Policies tab.

5 Click *Properties*.

The *Set Commands* tab displays.

6 Click *Add*.

The NetWare Server SET Command Wizard opens.

7 Select the server containing the SET parameters, then click *Next*.

IMPORTANT: The Policy/Package Agent must be running.

8 Select all of the commands you want to configure in the policy.

You can select whole categories by selecting the check box for the category, or clicking the plus sign to expand a SET command category and selecting the check boxes for individual commands to be included.

WARNING: Do not select the Set Developer Option SET command and change the default of Off to On. This parameter is meant to help developers debug server abends. It disables some of the operating system checking to prevent certain abends from occurring. Also, if the Set Developer Option is enabled, running NCP™ scripts that require keyboard entry could abend the server.

9 Click *Finish* when you are finished selecting the commands.

The selected commands are now displayed in the Set Commands tab for the policy.

10 To edit a SET command, click its plus sign to expand its attributes.

11 To edit an attribute, select the attribute, then click *Edit*.

A dialog box is displayed in which you can make changes to the attribute.

12 Repeat **Step 11** for each attribute to edit for a given SET command.

13 Repeat **Step 10** through **Step 12** to edit another SET command's attributes.

14 Schedule the policy (see **Section 4.7, "Scheduling Policies," on page 234**).

15 Click *OK* to close the policy.

If you click *Cancel*, neither the schedule or the SET parameter changes are saved.

Prohibited File

This policy allows you to monitor and enforce the deletion or moving of unauthorized files from a specified volume/drive or directory/folder. For example, you can automate deletion of .jpg, .mp3, and .avi files from a server.

All platforms are supported (NetWare®, Windows, Linux, and Solaris), including the use of the General tab.

With this policy, you can:

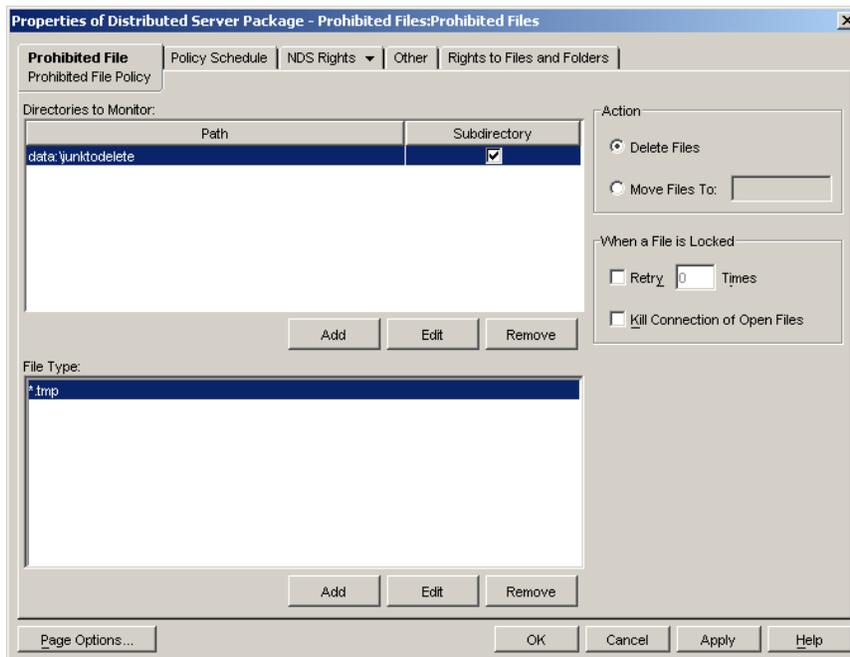
- ◆ Specify one or more volumes/drives or directories to monitor. You have the option to include all subdirectories.
- ◆ Specify which file types to monitor using wildcard combinations.
- ◆ Specify the action for all encountered files as follows:
 - ◆ Delete

- ◆ Move to specified location
- ◆ Specify a schedule for enforcement of the policy.

To configure a policy to manage prohibited files:

- 1 In ConsoleOne, right-click the *Distributed Server Package*, then click *Properties*.
- 2 Click *Policies > NetWare* (or other platform).
- 3 Click *Add*, then select *Prohibited File*.
- 4 Provide a unique name for the policy, then click *OK*.

The following property page is displayed:



5 Fill in the fields:

Directories to monitor: For this instance of the policy, you can specify the paths to be monitored:

- ◆ **Path:** This can be a volume, drive, or directory name. It must be the full path when a directory is given.

You can add multiple paths. For each path that you enter, files matching the file types that you define in the File Type field are either deleted or moved according to which Action button you select.

Variables are supported in the paths.

- ◆ **Subdirectory:** Select the check box to specify that all subdirectories be included.
If you want only a certain subdirectory, you should create another policy just for that subdirectory by giving its full path in the Path field. However, you cannot move files to a directory that is being monitored, or to any of its subdirectories.
- ◆ **Add:** Opens a dialog box where you can select a path. This field cannot be browsed, so you must know the full path to the files to be moved or deleted.
- ◆ **Edit:** Allows you to edit the selected path.

- ♦ **Remove:** Removes the selected path entry from the list.

Files to manage: You can specify the type of files you want to monitor:

- ♦ **Add:** Opens a dialog box where you specify a file type. You can use wildcards in the path:

* = any number of characters

? = any single character in that position

??? = any multiple characters in those positions

This field cannot be browsed, so you must specify the correct information to identify the files to be moved or deleted.

IMPORTANT: The ? wildcard acts differently in ZENworks than in DOS. For example, the search string *.htm? finds only files that end in .html, whereas DOS finds files that end in both .htm and .html. In other words, use of the ? wildcard in ZENworks means that you expect a character to occupy its position in the filename.

- ♦ **Edit:** Allows you to edit the selected file type.
- ♦ **Remove:** Removes the selected file type from the list.

Action: You have two options for how to handle the files you've specified in the Directories to Monitor and the Files to Manage boxes:

- ♦ **Delete files:** Select the option to delete the specified files from the locations you have identified.
- ♦ **Move files to:** Select the option to move the specified files to the path that you specify in this field. This field cannot be browsed, so you must know the full path to where you want the files to be moved.

If you move files:

- ♦ The full paths of the files are preserved (meaning if the path doesn't exist at the target, it is created there)
- ♦ Files are overwritten if they exist in the same path
- ♦ File or directory attributes and trustees are not transferred
- ♦ File ownership is preserved

IMPORTANT: If a directory is being monitored, you cannot move files into it or any of its subdirectories.

When a file is locked: Occasionally, files you might be trying to delete or move might be open. For these files, you can specify one of the following resolutions:

- ♦ **Retry ___ times:** Select the check box and enter a number for how many times you want to retry deleting or moving the file before continuing with the next file. Valid entries are from 1 to 10. The time used by each increment depends on the various hardware and software speeds involved in your system.

Use this field to allow enough time for a temporarily opened file to be closed, such as a file that is only opened long enough for the application to either obtain a copy for editing or write a new copy of the file.

- ♦ **Kill connection of open files:** (NetWare only) Kills the connection that is holding the file open so that the file can be deleted or moved, even if opened by a user at the time.

IMPORTANT: You can only kill connections to files on workstations. Server files cannot be disconnected from the process that has them open.

- 6 Click *OK* to close the policy.

Scheduled Down

You can automate when and how you want a server to go down, and whether it should be automatically brought back up.

To configure a scheduled downing for a server:

- 1 In ConsoleOne, right-click Distributed Server Package, then click *Properties*.
- 2 Click *Policies > NetWare* (or other platform).
- 3 Click *Add*, then select *Scheduled Down*.
- 4 Provide a unique name for the policy, then click *OK*.

Because the policies selected from this dialog box are plural, you can have multiple Scheduled Down policies listed on the *Policies* tab. Therefore, provide a unique name for this policy.

When you click *OK* after naming the Scheduled Down policy, the policy is selected on the *Policies* tab.

- 5 Click *Properties*.

The *Up Procedure* tab displays.

- 6 Select the downing method:

| Downing Option | Description |
|----------------|--|
| Reset Server | Downs the server and then does a warm boot |
| Restart Server | Downs the server and then restarts it |
| Down Server | Downs the server, does not restart it |

- 7 Schedule the policy (see [Section 4.7, “Scheduling Policies,” on page 234](#)).

- 8 Click *OK* to close the policy.

If you click *Cancel*, neither schedule for your newly scheduled Down policy is saved.

Scheduled Load/Unload

You can automate scheduled loading and unloading of NLM files and Java Class processes, and Linux and Solaris executables.

To configure the schedules:

- 1 In ConsoleOne, select the Distributed Server Package’s container, right-click the *Distributed Server Package*, then click *Properties*.
- 2 Click *Policies > NetWare* (or other platform).
- 3 Click *Add*, then select *Scheduled Load/Unload*.
- 4 Provide a name for this Load/Unload policy, then click *OK*.

Because the policies selected from this dialog box are plural, you can have multiple Load/Unload policies listed on the *Policies* tab. Therefore, provide a unique name for this policy. When you click *OK* after naming the Load/Unload policy, it is selected on the *Policies* tab.

5 Click *Properties*.

The *Scheduled Load/Unload* tab displays.

6 Click *Add*.

7 Select one of the following options:

[Section D.1, “Load NLM/Process,” on page 411](#)

[Section D.2, “Load Java Class,” on page 411](#)

[Section D.3, “Unload Process,” on page 412](#)

[Section D.4, “Start Service,” on page 412](#)

[Section D.5, “Stop Service,” on page 412](#)

Select an item for further instructions on configuring it.

8 Repeat **Step 6** and **Step 7** for each NLM or process to be included.

9 To rearrange the order, use the arrow keys.

10 Schedule the policy (see [Section 4.7, “Scheduling Policies,” on page 234](#)).

11 Click *OK* to close the policy.

If you click *Cancel*, your newly scheduled Load/Unload policy is not saved.

Server Down Process

You can automate the procedures your servers use when they are downed.

IMPORTANT: For the Windows, Linux, and Solaris platforms, if you down the server from its console, this policy is not recognized. Instead, you must down the server using the *Actions* option in *Remote Web Console* in iManager so that this policy can be applied.

To configure the downing process for a server:

1 In ConsoleOne, select the Distributed Server Package’s container, right-click the *Distributed Server Package*, then click *Properties*.

2 Click *Policies > NetWare* (or other platform).

3 Select the Server Down Process policy, then click *Properties*.

If the box under the *Enabled* column is not selected for the Server Down Process policy, select it before clicking *Properties*. A policy must be enabled to activate the *Properties* button.

4 To configure procedures for downing, select the *Down Procedure* tab, then click *Down Procedures*.

5 To enable the policy’s options, select the check box labeled *Follow this procedure when a down server is triggered*, then enter the number of minutes to wait before downing the server.

6 To disable login before downing, select the check box, then enter the number of minutes before downing to disable login.

7 To drop connections before downing, select the check box, then enter the number of minutes before downing the server to drop connections.

- 8** To configure an order for unloading, select the *Down Procedure* tab, then click *Ordered Unload*.
 - 8a** To include NLM files and processes, select the *Unload these NLMs and kill these processes in this order before downing* check box.
 - 8b** Click *Add*.
 - 8c** Select either *NLM* or *Process*, provide the name, then click *OK*.
 - 8d** To change the order, use the arrow keys.
- 9** To configure reporting, select the *Notification* tab, then click *Reporting*.
 - 9a** To have another server send an SNMP alert if the server is not up after a specified time, select the *Send SNMP Alert* check box, then enter the number of minutes.

For information about displaying SNMP traps on your management console, see [Section 4.3.1, “Compiling Zentrap.mib,” on page 206](#).
 - 9b** To specify which servers can watch for the restart and send the alert in case of failure, click *Add* to display an ordered list of candidate servers.

Policy and Distribution Services starts at the top of the list to communicate with the first server and use it for the alert notification. If Policy and Distribution Services cannot communicate with a server, the next one on the list is tried. The first server that can be used is the one that is scheduled to send the alert.
 - 9c** Browse to select a server.
 - 9d** Repeat [Step 9a](#) through [Step 9c](#) for each server needed.
 - 9e** To change the order, use the arrow keys.
- 10** To configure broadcast messages, select the *Notification* tab, click *Broadcast Messages*, then click *Send messages to connected users*.
 - 10a** Enter the number of times to send the message.
 - 10b** To broadcast custom text, enter it in the box.
 - 10c** To include the predefined message containing a time as the last line of your broadcast, select the check box.

The *x* minutes is derived from dividing the number of times from [Step 10a](#) into the number of minutes remaining before the server can be downed, then subtracting that amount (in whole minutes) for the amount to display in each broadcast. For example, if there are 10 minutes remaining and you select 5 in [Step 10a](#), the message is broadcast every two minutes. The number of minutes remaining after each broadcast will be two minutes less than at the last broadcast.
- 11** To configure targeted messages, select the *Notification* tab, click *Targeted Messages*, then click *Send e-mail to selected users when server is going down*.
 - 11a** To specify the users, groups, or e-mail addresses to receive the targeted messages, click *Add*.
 - 11b** Select either *User*, *Group*, or *E-Mail Address*.
 - 11c** Browse to select the user or group, or provide the e-mail address.
 - 11d** Repeat [Step 11a](#) through [Step 11c](#) for other users, groups, or e-mail addresses.
- 12** To configure the conditions for downing a server, select the *Conditions* tab, then click *Use Conditions*.
 - 12a** To specify the conditions, click *Add*.

12b Select from the following conditions to specify when not to bring the server down:

Some of these conditions require you to enter valid names. Others use the Select Object dialog box to browse for them.

File open: If the files that you specified are open. For example, a `.exe`.

NLM loaded: If the NLM files that you specified are running.

Server connected: If the server that you specified is connected.

User connected: If the users that you specified are connected.

Number of user connections: If the number of users connected exceeds the number you specify. In other words, don't bring the server down if too many users would be affected.

Workstation connected: If the workstations that you specified are connected.

12c Repeat **Step 12a** and **Step 12b** for each condition to add to the list.

12d To change the order, use the arrow keys.

13 Click *OK* to close the policy.

If you click *Cancel*, none of the Server Down Process policy changes made on any of the tabs are saved.

Server Scripts

You can automate script usage by your NetWare servers.

To configure server scripts:

1 In ConsoleOne, select the Distributed Server Package's container, right-click the *Distributed Server Package*, then click *Properties*.

2 Click *Policies > NetWare* (or other platform).

3 Click *Add*, then select *Server Scripts*.

4 Provide a unique name for the policy.

Because the policies selected from this dialog box are plural, you can have multiple Script policies listed on the *Policies* tab. Therefore, provide a unique name for this policy.

When you click *OK* after naming the Script policy, it is selected on the *Policies* tab.

5 Click *Properties*.

The *Script* tab displays.

6 Click *Add*, then select *Server Scripts*.

7 Provide a script name.

Script #1 displays.

8 Select the script type (NCF, NetBasic*, PERL).

IMPORTANT: NetBasic is not supported on NetWare 6.5 servers.

9 Enter the script text.

10 Repeat **Step 6** through **Step 9** for each script to be added.

11 Use the arrow keys to arrange the order to execute the scripts.

12 Schedule the policy (see **Section 4.7, "Scheduling Policies,"** on page 234).

- 13 Click *OK* to close the policy.

If you click *Cancel*, neither the schedule or any of the scripts entered are saved.

SMTP Host

You can set the TCP/IP address of the relay host that processes outbound Internet e-mail.

To configure the SMTP Host policy:

- 1 In ConsoleOne, right-click the *Service Location Package*, then click *Properties*.
- 2 Select the SMTP Host policy, then click *Properties*.

If the box under the *Enabled* column is not selected for the SMTP Host policy, select it before clicking *Properties*. A policy must be enabled to activate the *Properties* button.

The *SMTP Host* tab defaults.

- 3 Provide the TCP/IP address or DNS name (such as `mail.novell.com`), then click *OK* to close the policy.

If you click *Cancel*, the TCP/IP address is not saved.

SNMP Community Strings

This policy provides configuration and scheduling of SNMP community strings.

Make sure that you have compiled `zentrapp.mib` (see [Section 4.3.1, “Compiling Zentrapp.mib,” on page 206](#)).

IMPORTANT: Running `INETCFG` does not show that the policy has been applied to the server. Instead, use `TCPCON` to verify. See [“Verifying Community String Changes” on page 225](#).

To configure the SNMP Community Strings policy:

- 1 In ConsoleOne, select the Distributed Server Package’s container, right-click the *Distributed Server Package*, then click *Properties*.
- 2 Click *Policies > NetWare* (or other platform).
- 3 Select the SNMP Community Strings policy, then click *Properties*.

If the box under the *Enabled* column is not selected for the SNMP Community Strings policy, select it before clicking *Properties*. A policy must be enabled to activate the *Properties* button.

The *SNMP Community Policy* tab displays.

- 4 Fill in the *Community Strings* fields:

Monitor
Control
Trap

Community strings are case sensitive. Enter a string for each field as needed.

- 5 Select the *Schedule* tab, then schedule the policy (see [Section 4.7, “Scheduling Policies,” on page 234](#)).
- 6 Click *OK* to close the policy.

Verifying Community String Changes

To confirm that the SNMP Community Strings policy has been successfully applied to a server, do the following on any NetWare server:

- 1 At the server's main command prompt, enter `tcpcon` to display the following menu:

```
TCP/IP Console 6.00k                               NetWare Loadable Module

Host: Local System
Uptime: 0 Days 0 Hours 9 Minutes 42 Seconds
System: Novell NetWare 5.60.04 December 12, 2003

IP Received: 915      TCP Received: 9,222
IP Sent: 391         TCP Sent: 9,250
IP Forwarded: DISABLED  TCP Connections: 52

Available Options
SNMP Access Configuration
Protocol Information
IP Routing Table
Statistics
Interfaces
Display Local Traps

View and change the TCPCON options.
ENTER>Select ESC=Exit Menu                               F1=Help
```

- 2 Select *SNMP Access Configuration* to display “Local System” in the *Transport Protocol* field:

```
TCP/IP Console 6.00k                               NetWare Loadable Module

Host: Local System
Uptime: 0 Days 0 Hours 11 Minutes 33 Seconds
System: Novell NetWare 5.60.04 December 12, 2003

IP Received: 1,108   TCP Received: 10,108
IP Sent: 512        TCP Sent: 10,139
IP Forwarded: DISABLED  TCP Connections: 51

Available Options
SNMP Access Configuration
Transport Protocol: Local System
Host:
Community Name: public
Timeout: 5 (seconds)
Poll Interval: 1 (seconds)

The transport protocol for remote SNMP access.
ENTER>Select ESC=Previous Menu                               F1=Help
```

- Press Enter to display the *Transport* options:

```

TCP/IP Console  6.00k                               NetWare Loadable Module
-----
Host:  Local System
Uptime:  0 Days  0 Hours 11 Minutes 59 Seconds
System:  Novell NetWare 5.60.04  December 12, 2003

IP Received:  1,159      TCP Received:   10,276
IP Sent:      543        TCP Sent:      10,308
IP Forwarded: DISABLED   TCP Connections: 51

SNMP Acc
Transport Protocol: Local Syst
Host:
Community Name:   public
Timeout:          5 (seconds)
Poll Interval:   1 (seconds)

The transport protocol for remote SNMP access.
ENTER=Select ESC=Previous Menu                               F1=Help
  
```

- Select the *TCP/IP* option to display the TCP/IP transport protocol information:

```

TCP/IP Console  6.00k                               NetWare Loadable Module
-----
Host:  Local System
Uptime:  0 Days  0 Hours 12 Minutes 17 Seconds
System:  Novell NetWare 5.60.04  December 12, 2003

IP Received:  1,237      TCP Received:   10,437
IP Sent:      563        TCP Sent:      10,469
IP Forwarded: DISABLED   TCP Connections: 53

Available Options

SNMP Access Configuration
Transport Protocol: TCP/IP
Host: lb
Community Name:   public
Timeout:          5 (seconds)
Poll Interval:   5 (seconds)

The name or address of the host. DEL=Local System INS=Display host names.
ENTER=Select ESC=Previous Menu                               F1=Help
  
```

- Replace `1b` with the IP address of the NetWare server where you want to verify the string changes, and replace `public` with a valid monitor read string:

```

TCP/IP Console 6.00k                               NetWare Loadable Module
-----
Host: Local System
Uptime: 0 Days 0 Hours 13 Minutes 44 Seconds
System: Novell NetWare 5.60.04 December 12, 2003
-----
IP Received: 1,430      TCP Received: 11,118
IP Sent: 690           TCP Sent: 11,157
IP Forwarded: DISABLED TCP Connections: 51
-----
Available Options
-----
SNMP Access Configuration
-----
Transport Protocol: TCP/IP
Host: 192.68.1.203
Community Name: myreadstring
Timeout: 5 (seconds)
Poll Interval: 5 (seconds)
-----
The timeout interval for the request reply.
ENTER=Select ESC=Previous Menu                               F1=Help

```

- Press Esc to display the *Save TCP/IP Console Option?* menu, then select *Yes* to continue:

```

TCP/IP Console 6.00k                               NetWare Loadable Module
-----
Host: Local System
Uptime: 0 Days 0 Hours 14 Minutes 14 Seconds
System: Novell NetWare 5.60.04 December 12, 2003
-----
IP Received: 1,475      TCP Received: 11,298
IP Sent: 722           TCP Sent: 11,337
IP Forwarded: DISABLED TCP Connections: 51
-----
Available Options
-----
SNMP Access Configuration
-----
Transport Protocol
Host: Save TCP/IP Console Options?
Community Name:
Timeout: No
Poll Interval: Yes
-----
The timeout interval for the request reply.
ENTER=Select ESC=Previous Menu                               F1=Help

```

- 7 At this point, you should see the statistics being updated; however, if the community string changes are not displayed (as depicted below), make sure that the correct monitor string was entered in [Step 5](#).

```

TCP/IP Console  6.00k                               NetWare Loadable Module
-----
Host:  192.68.1.203
Uptime: Unknown
System: Unknown
-----
IP Received:  0          TCP Received:  0
IP Sent:      0          TCP Sent:      0
IP Forwarded: Unknown   TCP Connections: 0
-----
Available Options
-----
SNMP Access Configuration
Protocol Information
IP Routing Table
Statistics
Interfaces
Display Local Traps
-----
View and change the TCPCON options.
ENTER=Select ESC=Exit Menu          F1=Help

```

- 8 Another way to see that the policy is actually applied when the policy is deployed is to change the messaging level for the server’s Subscriber object to Level 4 or Level 5 (see the *SNMP trap* field in [Step 3](#) under [Section 3.6.3, “Configuring Subscribers,”](#) on page 150), then view the new and old string values in the TCP/IP Console screen as the changes occur.

SNMP Trap Targets

You can set targets for SNMP traps for the Policy/Package Agent.

- ♦ [“Understanding How the Windows Trap Target Policy Enforcer Behaves”](#) on page 228
- ♦ [“Configuring the SNMP Trap Target Policy”](#) on page 229

For information about displaying SNMP traps on your management console, see [Section 4.3.1, “Compiling Zentrap.mib,”](#) on page 206.

Understanding How the Windows Trap Target Policy Enforcer Behaves

The following abbreviations are used in this section to represent these Windows registry locations:

- ♦ **AGENT_KEY:**

```
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\SNMP\Parameters
```

- ♦ **ZFS_KEY:** HKEY_LOCAL_MACHINE\SOFTWARE\Novell\Zenworks\Zfs

The Windows SNMP trap target policy enforcer performs in the following sequence:

1. The policy enforcer first verifies an installation of an SNMP agent. This is done by checking if AGENT_KEY exists. If it exists, the enforcer assumes that an SNMP agent is installed and continues with the following steps. Otherwise, an error is returned and the processing stops.
2. The enforcer keeps track of all trap targets added by the ZENworks Server Management policy by placing the trap targets in ZFS_KEY. The trap targets are organized like the trap targets in

AGENT_KEY with a subkey of TrapConfiguration. The subkey TrapConfiguration contains community strings that are represented as registry subkeys. These community strings contain the trap target values associated with each community string.

3. Each trap target in the ZENworks Server Management policy is put into AGENT_KEY, unless it already exists. The policy enforcer ensures that each Server Management trap target is found, or is added to each community string. If no community strings exist in AGENT_KEY, a community string named “public” is created.
4. Any previously added trap targets found in ZFS_KEY that are removed from the ZENworks Server Management policy are removed from AGENT_KEY. Trap targets not added by Server Management are not removed.
5. If Microsoft’s SNMP agent is installed, the agent’s trap targets are automatically updated with registry changes.

Configuring the SNMP Trap Target Policy

To configure the SNMP Trap Targets policy:

- 1 In ConsoleOne, right-click the Service Location Package, then click *Properties*.
- 2 Select the SNMP Trap Targets policy, then click *Properties*.

If the box under the *Enabled* column is not selected for the SNMP Trap Targets policy, select it before clicking *Properties*. A policy must be enabled to activate the *Properties* button.

- 3 Click *Add*.
- 4 Provide a new target, then click *OK*.

TIP: Provide the TCP/IP address or DNS name of the target server. IPX addresses are not supported.

- 5 Repeat **Step 3** through **Step 4** for each new trap target.
- 6 Select the *Schedule* tab, then schedule the policy (see **Section 4.7, “Scheduling Policies,” on page 234**).
- 7 Click *OK* to close the policy.

If you click *Cancel*, none of the targets that you provided are saved.

Text File Changes

You can automate changes to text files on your servers.

To configure text file changes:

- 1 In ConsoleOne, select the Distributed Server Package’s container, right-click the *Distributed Server Package*, then click *Properties*.
- 2 Click *Policies > NetWare* (or other platform).
- 3 Click *Add*, then select *Text File Changes*.
- 4 Provide a unique name for the policy.

Because the policies selected from this dialog box are plural, you can have multiple text file policies listed on the *Policies* tab. Therefore, provide a unique name for this policy.

When you click *OK* after naming the text file policy, it is selected on the *Policies* tab.

- 5 Click *Properties*.

The *Text Files* tab defaults.

6 Click *Add*.

After one text file has been added, you are given the opportunity to select whether you are adding another text file or another change item for the selected text file.

To add another text file, select *Text File*. It does not matter which text file or change item is selected in the left pane—the text file is added to the far left level.

To add another change to a text file, in the left pane select the text file for the change, click *Add*, then select *Change*. The change item is added under the selected text file.

7 If you are adding a text file, provide the name of the text file.

8 Accept the default name (such as Change #1) or rename it; if you are adding a text file, click *OK*.

9 Click the down-arrow for the *Change Mode* field, then select the change mode from the drop-down list.

10 Click the down-arrow for the *Search Type* field, then select the search type from the drop-down list.

11 Enter the exact search string.

12 Select the check box if you want the string search to be case sensitive.

13 To find all occurrences of the search string, make sure the box is selected, or deselect the box to find only the first occurrence.

14 Click the down-arrow for the *Result Action* field, then select the action from the drop-down list that should result if a string is matched.

15 If you are replacing a string or entering a new one, enter the text in the *New String* text box.

16 Repeat **Step 6** through **Step 15** for each text file to add or each change to be made.

17 To reorder the text files and change items, use the arrow keys.

18 Schedule the policy (see **Section 4.7, “Scheduling Policies,” on page 234**).

19 Click *OK* to close the policy.

If you click *Cancel*, neither the schedule or any of the text files entered are saved.

ZENworks Database

If you installed the Server Management database during installation, but the database file is not associated with a Database object, you can set its object's DN so that the server this policy is associated with can find the database file for logging information.

To configure the ZENworks Database policy:

1 In ConsoleOne, right-click the *Service Location Package*, then click *Properties*.

2 Select the ZENworks Database policy, then click *Properties*.

If the box under the *Enabled* column is not selected for the ZENworks Database policy, select it before clicking *Properties*. A policy must be enabled to activate the *Properties* button.

3 Select the *Policy/Distribution Management* tab.

The *Inventory Management* tab defaults. Make sure you are using the correct tab.

4 Provide the DN of your ZENworks Database object, or browse to select the DN, then click *OK* to close the policy.

If you click *Cancel*, the DN is not saved.

ZENworks Server Management

This policy provides basic configuration parameters for Policy and Distribution Services.

To configure the ZENworks Server Management policy:

- 1 In ConsoleOne, select the Distributed Server Package's container, right-click the *Distributed Server Package*, then click *Properties*.
- 2 Click *Policies > NetWare* (or other platform).
- 3 Select the ZENworks Server Management policy, then click *Properties*.

If the box under the *Enabled* column is not selected for the ZENworks Server Management policy, select it before clicking *Properties*. A policy must be enabled to activate the *Properties* button.

The *General – Status* tab displays.

- 4 To determine the policy's general status:
 - 4a Select the procedure to follow when displaying messages at the server console.
 - 4b Select the procedure to follow when sending SNMP traps.

For information about displaying SNMP traps on your management console, see [Section 4.3.1, "Compiling Zentrap.mib," on page 206](#).
 - 4c Select the procedure to follow when recording information to a log file.

| Logging Procedure | Description |
|-------------------------------------|---|
| Log File | Select this option to enable it and provide the log file's filename. Include its full path. By default, Policy and Distribution Services uses <code>\zenworks\zfs-startup.log</code> , unless you enter a filename here. Then, for the servers where this policy is enforced, the log file you specify here is used instead of <code>zfs-startup.log</code> . Some examples: <code>sys:\zenworks\polpack.log</code> <code>sys:\zenworks\polpack.txt</code> <code>data:\zenworks\policies.log</code> |
| Delete Log Entries Older Than__Days | Use this option to control disk space usage. |
| E-Mail Messages | Select whether to send e-mail messages. The None or Errors Only options are recommended. |
| ◆ Users | You can add users, groups, or e-mail addresses. |
| ◆ Address Attribute | After you select users or groups, this field displays the attribute of the associated user or group. You can change the attribute from the drop-down list. |

IMPORTANT: Set the E-Mail Messages option to either None or Errors Only. If you set this to a more detailed level, performance degrades because of the extra e-mail messages that are created.

5 To determine the policy's configuration, select the *ZENworks Server Management* tab, then click *Configuration*.

5a Provide a console prompt.

You can customize the prompt using plain text and variables. The default is:

```
%SERVER_DN% - ZENworks Server Management >
```

You can use any of the predefined or user-defined variables (for more information, see [Section 9.2, "Types of Variables," on page 344](#)).

5b Provide a working path.

This is for Policy and Distribution Services temporary and backup files. The default directory is `\zenworks\pds\smanager\working`.

5c To determine how old database information should be before purging, enter the number of days.

All policy-related information older than the number of days entered is purged when Server Management is started on the same server where `zfslog.db` resides.

IMPORTANT: The database can only be purged if Server Management is running on the same server where `zfslog.db` is located.

Tiered Electronic Distribution information is purged manually from the database. For more information, see [Section 10.5, "Purging the Database," on page 359](#).

6 To set a port number for the ZENworks Web Server, select the *Port Configuration* tab and select or enter a port number.

7 Click *OK* to close the policy.

If you click *Cancel*, none of the policy changes on any of the tabs are saved.

4.3.5 Creating Custom Log Files Using Policies

If you want to create custom log files, you can use either the Tiered Electronic Distribution policy (Service Location Package) or the ZENworks Server Management policy (Distributed Server Package):

- ♦ **Tiered Electronic Distribution policy:** With this policy, you associate its Service Location Package to an eDirectory container, and all Distributor and Subscriber objects under it can use this policy. The Use Policy check box that is displayed in each of the object's properties allows you to individually select whether that Distributor or Subscriber should use the policy. The check box is disabled by default.

Using this policy, the Distribution Agent logs Tiered Electronic Distribution information to your custom log file for the selected Distributors and Subscribers.

- ♦ **ZENworks Server Management policy:** With this policy, you distribute its Distributed Server Package to the servers where you want the policy enforced.

Using this policy, the Policy/Package Agents for these servers log policy and software package information to your custom log file.

When you are creating and configuring one of these policies, the Path and Filename field for the log file is blank by default.

For information on how to create and configure these policies, see:

- ♦ [“Tiered Electronic Distribution” on page 209](#)
- ♦ [“ZENworks Server Management” on page 231](#)

4.4 Enabling Policies

A policy must be enabled before it is in effect for the policy package. You can disable a policy without removing it from the package.

To enable a policy:

- 1 In ConsoleOne, right-click the Policy Package object containing the policy to be enabled, then click *Properties*.
- 2 To enable a policy, select its check box under the *Enabled* column.
If you enable a policy, make sure it is correctly configured.
- 3 To cause an enabled policy to be enforced, distribute the policy package.
For more information, see [Section 4.5, “Distributing Policies,” on page 233](#).

4.5 Distributing Policies

You must distribute a Distributed Server Package before its policies are in effect. When you do distribute the package, its enabled policies are only in effect for the server where it is distributed after the Subscriber has extracted the Distribution.

To distribute policies to a server:

1. Create a Distribution that is a Policy Package type.
2. Configure the policies in the policy package.
3. Select a Channel for the Policy Package Distribution.
4. Subscribe the Subscribers to the selected Channel.
5. Send the Distribution.

The Policy/Package Agent on the receiving server extracts the enabled policies and enforces them on the server.

For more information on creating Policy Package Distributions, see [“Creating and Configuring the Distribution” on page 57](#).

4.6 Associating Policies

Because Distributors do not receive policies through Distributions, the Distributor object needs to be associated with the Container Package object so that it can use the Search policy for how to read the eDirectory tree when the Distributor is refreshed.

The Distributor object also needs to be associated with the Service Location Package. This package contains the ZENworks Database policy, which enables the Distributor Agent to locate the database file for writing report information. It also contains other policies the Distributor uses (see [Section 4.3.3, “Configuring Service Location Package Policies,” on page 207](#)).

To associate policy packages with the Distributor object's container:

- ◆ [Section 4.6.1, “Associating a Policy Package to the Distributor Object,” on page 234](#)
- ◆ [Section 4.6.2, “Associating the Distributor Object to a Policy Package,” on page 234](#)

4.6.1 Associating a Policy Package to the Distributor Object

To associate a policy package to the Distributor object's container:

- 1 In ConsoleOne, right-click the policy package, then click *Properties*.
- 2 Select the *Associations* tab, then click *Add*.
- 3 Browse to select the container where the Distributor object resides (or any container above it), then click *OK*.

If you click *Cancel*, the association you made is not saved.

4.6.2 Associating the Distributor Object to a Policy Package

To associate the Distributor object's container with a policy package:

- 1 In ConsoleOne, right-click the container where the Distributor object resides (or any container above it), then click *Properties*.
- 2 Select the *ZENworks* tab, click *Associated Policy Packages*, then click *Add*.
- 3 Browse to select the policy package, then click *OK*.
If you click *Cancel*, the association you made is not saved.
- 4 Repeat [Step 2](#) and [Step 3](#) for additional policy packages to be associated with the Distributor object's container.

4.7 Scheduling Policies

For information, see [Section 8.3, “Scheduling and Server Policies,” on page 338](#).

4.8 Viewing Effective Policies

To view which ZENworks 7 Server Management policies are in effect for the current server object:

- 1 At the `ZENworks Server Management` prompt on the server, enter `Policy List`.
Displays the policies that are currently in effect for the server.

4.9 Changing Policy Enforcement

You might need to change or stop policy enforcement for a particular server or a group of servers.

You can change policy enforcement in several ways:

- ◆ [Section 4.9.1, “Modifying a Policy That Is Being Enforced,” on page 235](#)
- ◆ [Section 4.9.2, “Stopping a Specific Policy From Being Enforced,” on page 235](#)
- ◆ [Section 4.9.3, “Removing Policy Enforcement for a Specific Subscriber,” on page 235](#)

- ◆ [Section 4.9.4, “Stopping Enforcement of a Policy Package Distribution,” on page 236](#)

4.9.1 Modifying a Policy That Is Being Enforced

To change a policy that is being enforced:

- 1 In ConsoleOne, right-click the Distributed Server Package object containing the policy to be modified, then click *Properties*.
- 2 Modify the policy as needed, then click *OK* to exit the policy package properties.

The next time the Distribution containing this policy package is built, the following transpires:

1. A new version of the Distribution is created because it had changed.
2. The Policy Package Distribution is sent according to the Send schedule of the Channel.
3. The Subscribers subscribed to the Channel each receive and extract the Policy Package Distribution according to their extraction schedules.
4. The modified policy is enforced on the Subscribers where the Policy Package Distribution was extracted.

4.9.2 Stopping a Specific Policy From Being Enforced

To stop a specific policy from being enforced:

- 1 In ConsoleOne, right-click the Distributed Server Package object containing the policy to be stopped, then click *Properties*.
- 2 Select the policy to be stopped, then do one of the following:
 - 2a Select the check box under the Enabled column to disable the policy.
 - 2b Click *Remove* to remove the plural policy.

You can delete plural policies from a policy package because they were previously added using the Add button.

- 3 Click *OK* to save the change and exit the policy package properties.

The next time the Distribution containing this policy package is built, the following transpires:

1. A new version of the Distribution is created because it had changed.
2. The Policy Package Distribution is sent according to the Send schedule of the Channel.
3. The Subscribers subscribed to the Channel each receive and extract the Policy Package Distribution according to their extraction schedules.
4. The disabled/removed policy is no longer enforced on the Subscribers where the Policy Package Distribution was extracted.

4.9.3 Removing Policy Enforcement for a Specific Subscriber

If you want to stop a distributed policy from being enforced on a specific Subscriber server, rather than on all Subscribers receiving that Distribution, do the following:

- 1 In ConsoleOne, right-click the Subscriber object, then click *Properties*.

- 2 Select the *Channels* tab, select the Channel containing the policy to be removed from enforcement, click *Remove*, then click *OK*.
- 3 Click *OK* to close the Subscriber object's properties.
- 4 On the Subscriber server's file system, delete the following files:
 - ♦ The Distribution directory containing the policy's Distribution file
 - ♦ The related Policy file (.pol) from the \smanager\policy directory (which was created when the Policy Package Distribution was extracted)
- 5 Reset the Subscriber server to refresh its policy configuration.

The Subscriber no longer receives the Policy Package Distribution containing that policy, nor does it continue to enforce the policy previously distributed to the Subscriber.

4.9.4 Stopping Enforcement of a Policy Package Distribution

If you need to stop enforcement of a Policy Package Distribution for all of the Subscribers where it was distributed, you must follow certain steps. Because the policy package was distributed, each Subscriber that received the Distribution can still enforce that policy if you only delete the policy package object.

To stop enforcement, do the following:

- 1 In ConsoleOne, delete the Distribution object for the Policy Package type.

IMPORTANT: If the policy package has other policies that you do not want to stop, then do not delete the package. Instead, just disable the policy that you want to stop.

- 2 On the Subscriber server's file system, delete the .pol file that was created by the Policy Package Distribution.

The .pol file is located under the \zenworks\pds\smanager\policies directory.

- 3 Refresh the policies on each Subscriber.

You can do this from each Subscriber server's console using the Policy Refresh command, or from iManager using the Refresh option.

The policies in the Policy Package Distribution are no longer enforced on the Subscriber after its policies have been refreshed. The refresh process clears its memory of all policies, then reloads them from the Policy Package Distributions existing in its file system.

Server Software Packages

5

Novell® ZENworks® Server Management provides the Server Software Packages component for managing files and applications on your network. Using software packages, you can automate the installation and upgrading of software on your servers.

The real value in using software packages is to set up processes to be done on a server before and after installation of the package.

The following sections give you an understanding of how you can benefit from using the Server Software Packages component:

- ♦ [Section 5.1, “Software Management through Server Software Packages,” on page 237](#)
- ♦ [Section 5.2, “Understanding Server Software Packages,” on page 237](#)
- ♦ [Section 5.3, “Planning Server Software Packages,” on page 249](#)
- ♦ [Section 5.4, “Setting Up Server Software Packages,” on page 251](#)
- ♦ [Section 5.5, “Using Server Software Packages to Delete Directories on Servers,” on page 268](#)

5.1 Software Management through Server Software Packages

Software management is done by creating Server Software Packages and distributing them using Tiered Electronic Distribution. You can configure Server Software Packages so that a server must meet certain minimum requirements before a package is installed on it. Software packages can consist of multiple software package components.

Each software package component can also be configured so that minimum requirements must be met before that component can be installed on the server.

5.2 Understanding Server Software Packages

Policy and Distribution Services provides the means to automate and standardize the distribution and installation of server files and applications. This includes your ability to standardize NLM™ versions, configuration files, databases, and more. Review the following sections:

- ♦ [Section 5.2.1, “Understanding Server Software Packages and Components,” on page 238](#)
- ♦ [Section 5.2.2, “Understanding Software Package and Component Configurations,” on page 238](#)
- ♦ [Section 5.2.3, “Determining the Installation Order of Software Packages,” on page 239](#)
- ♦ [Section 5.2.4, “Executing Extracted Files,” on page 240](#)
- ♦ [Section 5.2.5, “Compiling Software Packages,” on page 241](#)
- ♦ [Section 5.2.6, “Accessing Software Packages,” on page 241](#)
- ♦ [Section 5.2.7, “Distributing Software Packages,” on page 242](#)
- ♦ [Section 5.2.8, “Distributing Software Packages to a Cluster,” on page 242](#)
- ♦ [Section 5.2.9, “Managing Server Software Packages,” on page 243](#)

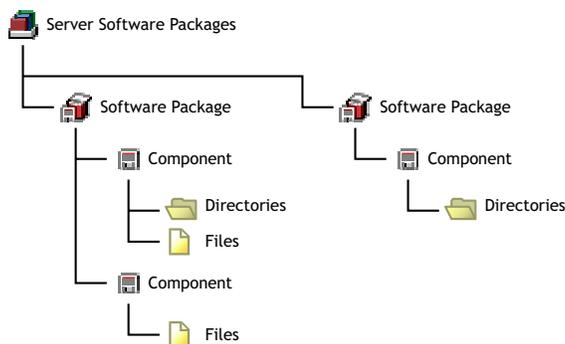
- ◆ [Section 5.2.10, “Failure of Software Package Installations,”](#) on page 248
- ◆ [Section 5.2.11, “Rolling Back Software Package Installations,”](#) on page 249

5.2.1 Understanding Server Software Packages and Components

To distribute server files and applications for installation on a server, you must include the software in a software package. You create the software packages under the Server Software Packages namespace in ConsoleOne®. Creating software packages is like building a software installation executable.

Figure 5-1 illustrates the relationship between software packages and package components:

Figure 5-1 Server Software Packages Sample Tree



Note the following:

- ◆ Software Package objects are displayed under the Server Software Packages namespace
- ◆ A Software Package object can contain multiple Component objects
- ◆ Component objects can contain files and directories
- ◆ Each software package can include all of the files for one or several applications
- ◆ Software Package configuration files (.spk and .cpk) are stored on a server or workstation file system

5.2.2 Understanding Software Package and Component Configurations

Software packages and their components contain configuration information and installation requirements. Because each Component object is governed by its own set of configuration parameters and installation requirements, you might have multiple components for a software package, such as pre-installation actions, installation actions, and post-installation actions.

You can configure every aspect of the distribution and installation of server files and applications, including the following:

- ◆ Requiring a specific operating system
- ◆ Specifying how much RAM the target server needs
- ◆ Specifying how much disk space the target server needs

- ◆ Requiring certain SET commands on the target server
- ◆ Making changes to the target server's registry
- ◆ Replacing files on the target server
- ◆ Requiring specific `products.dat` entries

Requiring Software Package Installation Prerequisites

Not only can a software package have installation prerequisites, but each of its components can also have its own installation prerequisites. The hierarchy for adhering to prerequisites to determine installation eligibility is:

- ◆ If the prerequisites for the package are not met, none of the components are installed.
- ◆ If the prerequisites for the package are met, the components are eligible to be installed.
- ◆ If the prerequisites for a component are not met, that component is not installed.

Because some components can be installed while others are not, a partial installation of the software package is possible.

IMPORTANT: When you specify prerequisites, be sure to create prerequisites at the software package level that would apply equally to all of its components, and create prerequisites at the component level that are specific to that component.

Naming Software Packages

When you create a software package, you initially give it a `.spk` extension, which represents a software package that has not yet been compiled. This file contains all of the installation requirements for the software package and all of its components.

WARNING: Do not use double-byte characters in the software package name. This causes an error in any report you run on the software package.

5.2.3 Determining the Installation Order of Software Packages

There are two issues concerning the ordering of Server Software Packages in Distributions:

- ◆ [“Forcing the Software Package Distribution Order Using Multiple Distributions” on page 239](#)
- ◆ [“Forcing the Software Package Distribution Order Using Dependencies” on page 240](#)
- ◆ [“How Rollback Is Affected by Software Package Ordering” on page 240](#)

Forcing the Software Package Distribution Order Using Multiple Distributions

If you want to include multiple software packages in one Distribution, consider the following:

- ◆ Multiple software packages are not gathered into a Distribution in any particular order when the Distribution is built
- ◆ Multiple software packages are not applied to a server in any guaranteed order when the Distribution is extracted and installed
- ◆ Multiple software packages that are contained in one Distribution and start their installations in a certain order might not all finish in that same order

To install software packages in a particular order:

- 1 Place each software package in its own Distribution (one software package per Distribution).
- 2 Control the order of software package installations by scheduling the order when the Distributions are sent and extracted.

Forcing the Software Package Distribution Order Using Dependencies

Another way to ensure software package installation order is to use dependencies with multiple Software Package Distributions, such as placing a dependency in a subsequent software package that is established in previous software package.

For example:

1. Create Software Package Distribution 1.
2. Create Software Package Distribution 2 with a dependency on something that is installed from Software Package Distribution 1.
3. Send both Distributions.
4. If the Subscriber attempts to extract Software Package Distribution 2 first, it will fail.
5. The Subscriber extracts Software Package Distribution 1, which provides the dependency on the Subscriber required by Software Package Distribution 2.
6. The Subscriber can now successfully extract Software Package Distribution 2.

With this scenario, you do not need to use schedules to control the installation order.

How Rollback Is Affected by Software Package Ordering

Rollback is also affected by the fact that multiple software packages contained in one Distribution won't necessarily finish extracting in the same order that they started.

Although you can specify the order for processing software packages that are contained in a Distribution, this order is not guaranteed. This is because the length of time it takes for software packages to finish processing can be different for each package, and it is the finishing time for a software package that determines its rollback order.

In other words, you can only roll back the last software package that was successfully processed, and then other software packages only in the reverse order of when they finished processing.

You can use the Package List command to view the order in which software packages finished processing. An asterisk marks the next package that is available for rollback.

For more information on rollback, see [Section 5.2.11, "Rolling Back Software Package Installations," on page 249](#).

5.2.4 Executing Extracted Files

In a Software Package Distribution, some of the files that the software package copies to a server might be executables that you want to have execute in conjunction with extracting the Software Package Distribution.

To run executable files that are delivered through a software package, configure the pre or post execution actions, including order of execution, of the files in the software package. Pre and post

actions are available when creating the Server Software Package and when creating the Software Package Distribution.

For more information, see [Section 3.4.6, “Pre and Post Processing for Distributions,”](#) on page 126, and [“Configuring the Software Package Components”](#) on page 256.

5.2.5 Compiling Software Packages

After you have defined your software packages, including configuring the components, you must compile the software package. This process compresses the files and applications and their configurations into one file for distribution.

The default extension for a compiled software package is `.cpk`. The compiled version contains all of the files necessary to install the files and applications that the software package represents.

IMPORTANT: If you provide the path and filename of the `.spk` when you are prompted for the compiled filename, the `.spk` is overwritten and can no longer be edited. Be sure to use the `.cpk` extension when naming the compiled version.

A `.cpk` file has the potential to be very large (hundreds of megabytes), because software packages can include many large files to be copied. Therefore, `.cpk` files should generally be stored on a server where you have sufficient free disk space.

However, software packages can perform simple functions, which would make the `.cpk` files’ sizes relatively small, so that you could store them on a workstation. For example, a software package could be configured to just delete directories on a file server (see [Section 5.5, “Using Server Software Packages to Delete Directories on Servers,”](#) on page 268).

When a rollback-enabled software package is successfully installed, a rollback package is created on the server. Processing this rollback package returns the server to its original state (before the package was installed). For more information, see [Section 5.2.11, “Rolling Back Software Package Installations,”](#) on page 249.

5.2.6 Accessing Software Packages

Where you save software packages (on workstations or on servers) depends on how you want to manage the software packages.

Because the Server Software Packages component uses a namespace in ConsoleOne, it enables you to have access to software packages from any workstation or server where you are running ConsoleOne.

However, you should be aware of the following issues:

- ◆ [“Running ConsoleOne from a Workstation”](#) on page 242
- ◆ [“Running ConsoleOne from a Server”](#) on page 242

For information on managing software packages from multiple workstations, see [Section 5.2.9, “Managing Server Software Packages,”](#) on page 243.

Running ConsoleOne from a Workstation

If you run ConsoleOne from a workstation and save a software package to that workstation, the package is not available in ConsoleOne to other workstations or servers running ConsoleOne.

Running ConsoleOne from a Server

You must have the same drive mapping to a server on different workstations if you run ConsoleOne from the server at those workstations. Otherwise, any software package you save to that server cannot be read at the different workstations.

For example, the following scenario illustrates when a package can be found:

1. You run ConsoleOne from Workstation A to access Server A.
2. Server A is mapped as drive S: for Workstation A.
3. You save `pkg_a.spk` to Server A.
4. You run ConsoleOne from Workstation B to access Server A.
5. Server A is also mapped as drive S: for Workstation B.
6. `pkg_a.spk` can be found because both workstations were mapped to drive S:.

The following scenario illustrates when a package cannot be found:

1. You run ConsoleOne from Workstation A to access Server A.
2. Server A is mapped as drive S: for Workstation A.
3. You save `pkg_a.spk` to Server A.
4. You run ConsoleOne from Workstation B to access Server A.
5. Server A is mapped as drive T: for Workstation B.
6. `pkg_a.spk` cannot be found because you are looking for the package on drive T: when it was previously saved to drive S:.

The only difference between the scenarios is the drive letter mappings to Server A for each workstation.

5.2.7 Distributing Software Packages

Distributions can include software packages, which are installed, or file groupings, which are extracted.

The Policy/Package Agent extracts or installs Software Package Distributions on the Subscriber server.

When software packages are created, they can contain system requirements that must be met before you install the package on the target Subscriber's server. If the Subscriber meets these requirements, the subscription schedule determines when the package is actually installed.

5.2.8 Distributing Software Packages to a Cluster

When you send a Distribution containing software packages to a cluster to update the sys: volume for each node, the only node in the cluster that receives it is the one that currently has the Subscriber software running.

Because the machines comprising the nodes in the cluster run the Subscriber software, only one node at a time in a cluster is actively running the Subscriber software.

Therefore, if you want to use a Software Package Distribution to update files on a sys: volume for each node in a cluster, you must do this manually by updating one node, bringing it down so that the next node in the failover sequence sees that the previous node has failed and start running the Subscriber software, then update that machine, bring it down, and so on, until all of the machines in the cluster have been updated. Then restart all of the downed servers in the cluster and the primary node's machine takes over again.

You can use a Software Package Distribution to update files on the cluster machine itself, such as Tiered Electronic Distribution .ncf files, because the Subscriber software is contained on the cluster machine's shared hard drive.

5.2.9 Managing Server Software Packages

The following sections explain where to store Server Software Package files, and how to manage them:

- ♦ [“Understanding Server Software Package Files” on page 243](#)
- ♦ [“Understanding Your Software Package Management Options” on page 244](#)
- ♦ [“Storing and Managing .Spk Files Using One Workstation” on page 244](#)
- ♦ [“Storing .Spk Files on a Network Server and Managing Them from Multiple Workstations” on page 244](#)
- ♦ [“Example: Using a Master Snapinprefs.ser File” on page 247](#)

Understanding Server Software Package Files

There are three file types associated with software packages:

- ♦ **Configuration file (.spk):** When you create a Server Software Package, you initially create a configuration file (.spk) for it. This file's configuration is created in the properties of the software package object in the Server Software Packages namespace in ConsoleOne.

A .spk file is generally small (around 100 KB). Therefore, it can generally be stored on the workstation running the instance of ConsoleOne that you are using to create and manage software packages.
- ♦ **Compiled file (.cpk):** When you compile a software package, a .cpk file is created from the .spk file's configuration information. This provides the content of the software package, such as files or functions. The .cpk file is used to install the software package's content on a server.

You should generally store .cpk files on a server where there is sufficient free disk space, because compiled software packages might contain many files. However, you can store small .cpk files that only contain functions on a workstation.
- ♦ **Preferences file (.ser):** The preferences file (snapinprefs.ser) is automatically created on the workstation being used to create a software package. It contains pointers to the .spk files for the software packages.

This preferences file allows you to see the software packages in the namespace in ConsoleOne. In other words, software packages are displayed in the Server Software Packages namespace for an instance of ConsoleOne only if the .spk file's path is listed in the preferences file located on the workstation running that instance of ConsoleOne.

When you create a new software package, you specify the local path for the `.spk` file. When you compile a software package, you specify the server's path for the `.cpk` file. After you exit ConsoleOne, any time you have created, deleted, or compiled a software package, the `.spk` file paths are logged to the `snapinprefs.ser` file.

The path to the `.cpk` file is also logged to the `snapinprefs.ser` file. The next time you compile the software package, the wizard displays the `.cpk` file's previous location so that you do not need to remember it each time you compile the package. However, you need to note where you store the `.cpk` files for when you want to distribute them using Tiered Electronic Distribution, because the `.cpk` files' locations are not stored in the software package's properties.

Understanding Your Software Package Management Options

If you are using only one specific workstation for viewing, creating, and managing all of your software package files, then you can store the `.spk` files on that workstation.

It is also possible to manage your software packages from multiple workstations. This requires that you centralize your `.spk` file storage to a network server. This method also requires the use of a master `snapinprefs.ser` file so that you can view all of your software packages from any workstation.

The next two sections explain these management options.

Storing and Managing .Spk Files Using One Workstation

If you use only one workstation for viewing, creating, and managing your software packages, you can store the `.spk` files on the workstation and the `.cpk` files on a server.

Whether you are running ConsoleOne from the workstation where it is installed or from a workstation that uses an installation of ConsoleOne on a network server, the `snapinprefs.ser` file is updated on the workstation being used to run ConsoleOne.

Storing .Spk Files on a Network Server and Managing Them from Multiple Workstations

If you want to use multiple workstations for viewing, creating, and managing the same set of software packages, you need to store all `.spk` files on a network server so that they can be accessed by each workstation.

You might also want to use different workstations for managing different sets of software packages. Any workstation used to create `.spk` files has a software package preferences file of its own created on the workstation used to manage the software packages.

You can manage all of your software packages from multiple workstations if you use a master copy method for the `snapinprefs.ser` file.

- ◆ [“Understanding the Software Package Preferences File” on page 245](#)
- ◆ [“Managing Software Packages from Multiple Workstations” on page 245](#)
- ◆ [“General Rules for Managing Software Packages from Multiple Workstations” on page 246](#)
- ◆ [“The Best Scenario for Using Multiple Workstations to Manage Software Packages” on page 247](#)

Understanding the Software Package Preferences File

When you create a Server Software Package object in ConsoleOne, a software package preferences file (`snapinprefs.ser`) is created in the following location on the workstation running ConsoleOne:

```
c:\documents and settings\user_ID\.consoleone (Windows 2000)
```

where `user_ID` is the user directory associated with how you are logged in, such as Administrator.

The full path and filename for a software package is drive-dependent. The `snapinprefs.ser` file contains the drive letter, path, and package name for each `.spk` created by the workstation.

The `snapinprefs.ser` file is unique for each workstation. It is the preferences file that is updated whenever you add or remove `.spk` files using that workstation. Therefore, if you use three different workstations to create `.spk` files, you have three different `snapinprefs.ser` files, each on its own workstation.

When you start ConsoleOne, it checks to see if a `snapinprefs.ser` file was created for that workstation by the instance of ConsoleOne being run on the workstation, and whether ConsoleOne is installed on that workstation or is being run on that workstation from an instance installed on a server. If the file does not exist, a `snapinprefs.ser` file is created when you exit ConsoleOne. If it exists, the `snapinprefs.ser` file is updated with the full paths to any new `.spk` files.

You can copy a `snapinprefs.ser` file from one workstation to another. However, after replacing a `snapinprefs.ser` file with a copy from another workstation, you need to restart ConsoleOne to see any change.

A software package can become unusable if you change drive mappings after creating the package, because the `snapinprefs.ser` file's location to the package is then different. However, if you use a UNC path, this is not an issue as long as the workstation has access to that UNC path.

If you replace the `snapinprefs.ser` file on a workstation, you need to manually insert any software packages missing from the newly copied `snapinprefs.ser` file. Otherwise, the software packages listed in the `snapinprefs.ser` file that was replaced would be inaccessible on the workstation.

Even if a workstation has never been used to create a software package, you can copy a `snapinprefs.ser` file from another workstation to the appropriate location (`c:\...\consoleone`). Then, when you start ConsoleOne, you can see all of the software packages that are listed in the `snapinprefs.ser` file that you copied.

For more information, see [“Example: Using a Master Snapinprefs.ser File” on page 247](#).

Managing Software Packages from Multiple Workstations

If you are using multiple workstations for creating, deleting, and compiling the same set of software package files, you should do the following:

1. Store the `.spk` files on one network server (usually the server where you are storing their corresponding `.cpk` files), so that the software packages can all be accessed from any workstation.
2. When mapping a workstation to the server where the `.spk` and `.cpk` files are stored, use the same drive letter for all workstations.

3. Create a master `snapinprefs.ser` file to use for keeping all workstations updated with their latest software package additions, deletions, and compilations (see “[Setting Up the Master Snapinprefs.ser File](#)” on page 251).
4. Create a batch file for starting and stopping ConsoleOne on a workstation (see “[Creating and Using the ConsoleOne Batch File](#)” on page 252). This batch file does two things:
 - ◆ Automatically upload the latest `snapinprefs.ser` file from the storage server to the workstation any time ConsoleOne is started on that workstation.
This allows you to see all software packages from the workstation where you started ConsoleOne.
 - ◆ Automatically download the revised `snapinprefs.ser` file from the workstation to the storage server when ConsoleOne is exited on that workstation.
This creates a new master copy of the `.ser` file containing the workstation’s latest software package additions.
5. Run the batch file from any workstation where you want to manage software packages (see “[Using the ConsoleOne Batch File](#)” on page 254).

General Rules for Managing Software Packages from Multiple Workstations

Using a master copy for the `snapinprefs.ser` file works only if you exit ConsoleOne on one workstation, then start it on another workstation. This sequential method does not work for concurrently running instances of ConsoleOne where each instance is updating its local `snapinprefs.ser` file. The instance of ConsoleOne that is exited last overwrites the master copy with its local `.ser` file.

IMPORTANT: Creating, deleting, or compiling software packages in ConsoleOne are the only functions that cause logging to the `snapinprefs.ser` file. Therefore, you can use ConsoleOne to manage software packages, such as viewing and editing properties, without starting ConsoleOne from the batch file. Just make sure that you do not add, delete, or compile any `.spk` files in ConsoleOne if you do not start ConsoleOne with the batch file.

To manage software packages using this master copy/single server/multiple workstation method, observe the following general rules:

- ◆ Always exit ConsoleOne after creating a new software package (`.spk` file) or compiling a new package (`.cpk` file). This causes the master `snapinprefs.ser` file to contain the newest software package links.
- ◆ Never have two or more workstations concurrently managing software packages. The batch file used to start ConsoleOne on these workstations could cause paths to any newly created software packages to be lost.
- ◆ Never use the batch file to start ConsoleOne when you do not intend to manage software packages. Instead, start ConsoleOne without using the batch file.

You need to do this because the batch file always overwrites the master copy on the software package storage server when ConsoleOne is exited (if ConsoleOne was started by the batch file). You could inadvertently overwrite the master `snapinprefs.ser` file and lose links to newly created software packages.

For example, on Workstation_A you run the batch file to start ConsoleOne, do administrative work other than software packages, for some reason go to Workstation_B where you decide to create a new software package (so you use the batch file again), exit ConsoleOne on

Workstation_B, then later exit ConsoleOne on Workstation_A. Your new software packages created on Workstation_B no longer have links to them in the master `snapinprefs.ser` file.

The Best Scenario for Using Multiple Workstations to Manage Software Packages

The best scenario is that you have one administrator who can use multiple workstations to manage your software packages. If you have multiple administrators, they need to coordinate so that they don't overwrite each other's latest software package additions and deletions in the master `snapinprefs.ser` file.

For more information, see [“Example: Using a Master Snapinprefs.ser File” on page 247](#).

Example: Using a Master Snapinprefs.ser File

Keeping the master copy on the server properly updated is a matter of timing. For example, in the following scenario, the first `snapinprefs.ser` file was initially created on Workstation A, then copied to the network server to be the master `snapinprefs.ser` file. Both workstations are using Windows 2000.

A batch file is used to start ConsoleOne for the purpose of controlling events before and after using ConsoleOne. For example:

1. Administrator A starts the batch file on Workstation A to begin ConsoleOne.
2. The batch file running on Workstation A identifies the storage server as being mapped to drive M: (or it maps drive M: to that server).
3. The batch file copies the master `snapinprefs.ser` file from the server at drive M: to the `c:\documents and settings\user_ID\.consoleone` directory on Workstation A.
4. Administrator A creates a new software package, naming it `ssp1.spk`.
5. Administrator B starts the batch file on Workstation B to begin ConsoleOne.
6. The batch file running on Workstation B identifies the storage server as being mapped to drive M: (or it maps drive M: to that server).
7. The batch file copies the master `snapinprefs.ser` file from the server at drive M: to the `c:\documents and settings\user_ID\.consoleone` directory on Workstation B.

This is the same version of `snapinprefs.ser` that Administrator A had copied to Workstation A, except that it hasn't been updated yet with Administrator A's addition of `ssp1.spk`.

8. Administrator B creates a new software package, naming it `ssp2.spk`.
9. Administrator B exits ConsoleOne, which updates `snapinprefs.ser` on Workstation B with the `ssp2.spk` path.
10. The batch file running on Workstation B updates the master `snapinprefs.ser` file on the network server at drive M: with the updated `snapinprefs.ser` file from Workstation B. This updated master `snapinprefs.ser` file now contains the location of `ssp2.spk`.
11. Administrator A exits ConsoleOne, which updates `snapinprefs.ser` on Workstation A with the `ssp1.spk` path.

12. The batch file running on Workstation A updates the master `snapinprefs.ser` file on the network server at drive M: with the updated `snapinprefs.ser` file from Workstation A. This updated master `snapinprefs.ser` file now contains the location of `ssp1.spk`. However, the location for `ssp2.spk` has been lost, because Workstation B's update of the master `snapinprefs.ser` file was overwritten by Workstation A's later update.

This scenario would cause Administrator B to lose access to `ssp2.spk`, because the master `snapinprefs.ser` file no longer contains a record of `ssp2.spk`'s location. It was replaced with Administrator A's `snapinprefs.ser` file containing only `ssp1.spk`'s location. However, you can manually insert `ssp2.spk` into ConsoleOne (using the Insert Software Package option), so that it is listed in the `snapinprefs.ser` file along with `ssp1.spk`.

For this multiple-workstation management method to work, you must ensure that the master `snapinprefs.ser` file you keep on the network server is only used by one workstation at a time for creating, deleting, or compiling `.spk` files. However, you can use multiple workstations to simultaneously view or edit a Server Software Package object's properties, because the viewing and editing functions do not cause updates to a `snapinprefs.ser` file.

WARNING: You can perform edits to the properties of the Server Software Package object without affecting the `snapinprefs.ser` file. However, because Server Software Package objects are not in eDirectory™, but only in a name space, the `.spk` files might not have file-locking protection, unless the server's operating system provides this functionality. Therefore, you should devise management controls to protect against overwriting `.spk` files when using multiple workstations to manage software packages.

5.2.10 Failure of Software Package Installations

If a server fails to meet any of the software package requirements, it is not installed:

- ♦ [“Failure During an Installation” on page 248](#)
- ♦ [“Failure of a Component” on page 248](#)

Failure During an Installation

The system tracks all changes made by the installation of a software package. Except as noted under [Section 5.2.11, “Rolling Back Software Package Installations,” on page 249](#), if a server meets the requirements and the installation begins, then a failure condition halts the installation prematurely, the installation program automatically returns the server to the state it was in before the installation began, undoing what had been done to that point.

Failure of a Component

If a server meets the software package requirements, and some of the components meet the installation requirements and some do not, the installation is completed except for the components where the requirements were not met. In this case, you would have a partial installation of the package.

You should organize your software packages and their components so that if this happens, it does not leave disconnected or incomplete files or applications on the target machine.

5.2.11 Rolling Back Software Package Installations

Software package rollback is enabled by default. You should not disable rollback, unless you know the installation never needs to be undone.

- ♦ [“Rollback Methods” on page 249](#)
- ♦ [“Rollback of Older Installations” on page 249](#)
- ♦ [“Rollback Exceptions” on page 249](#)

Rollback Methods

There are two ways you can roll back a software package installation:

- ♦ On the server containing the package to be rolled back, enter `package rollback` at the server's ZENworks Server Management console prompt.
- ♦ Use a Web browser to access the ZENworks Server Management role and select the rollback option. For more information, see [Chapter 2, “Novell iManager,” on page 63](#).

The software package is uninstalled, leaving the server as if it had never been installed, except for any changes that might have been made to the server in using the installed application.

Rollback works, even if some components have not been installed during a successful package installation, because the installation program tracks what was and wasn't installed by the software package.

Rollback of Older Installations

When you roll back a software package installation, it is the last software package installed on that server. If that's not the one you need to roll back, you must roll back each installation, beginning with the more recent installations first, until you have rolled back the desired package.

For example, you installed three software packages on a server (Package1, Package2, and Package3). Package1 was installed first and Package3 was installed last. If you want to roll back Package2, you must first roll back Package3. To do so, you need to enter `package rollback` at the server's ZENworks Server Management console prompt once for Package3, then again for Package2.

Rollback Exceptions

You can normally undo a successful software package installation by rolling it back. However, any software package installation that runs a program such as a NetBasic script, a Java Class, or an NLM that modifies the server cannot be rolled back successfully, because those programs or services might have launched other programs that made changes on the server, which cannot be tracked for rolling back.

5.3 Planning Server Software Packages

Review each of the following sections and take notes as instructed. This information will help you to configure your software packages and their components.

- ♦ [Section 5.3.1, “Which Files or Applications Do I Want to Distribute?,” on page 250](#)
- ♦ [Section 5.3.2, “What Software Package Components Are Needed?,” on page 250](#)

- ◆ [Section 5.3.3, “What Minimum Requirements Are Needed?”](#) on page 250

After planning your software package, continue with [Section 5.4, “Setting Up Server Software Packages,”](#) on page 251.

5.3.1 Which Files or Applications Do I Want to Distribute?

You can distribute software packages containing files and applications for servers, as well as software packages containing end-user applications for further distribution in ZENworks Desktop Management to workstations. For information on configuring a Desktop Application object, see [“Application Management”](#) in the *Novell ZENworks 7 Desktop Management Administration Guide*.

If you have ZENworks 7 Desktop Management installed, you can also distribute desktop applications using Tiered Electronic Distribution, instead of including them in software packages. For more information, see [Chapter 6, “Desktop Application Distribution,”](#) on page 273.

You can include a file or application in more than one software package. For instance, a word processor application could be included in a software package designed for a secretarial group and one designed for a financial group.

Where applicable, organize the files and applications into logical groups for inclusion in software packages.

Follow the steps under [Section 5.4.2, “Creating a Server Software Package,”](#) on page 255 and [Section 5.4.4, “Creating the Software Package Components,”](#) on page 256 and note the information you need to know for creating the software package and its components.

5.3.2 What Software Package Components Are Needed?

You can have one or more components in a software package. For example, if you create a software package for installing virus protection software, you might want one component to be the original virus protection program, and another component a current virus pattern update file.

Components in a software package can each have the same or different installation requirements. If you give the components different requirements, they might not all be installed together. You can save time and minimize error by giving all of the components the same requirements.

IMPORTANT: You should include in the same component the files and applications that are dependent on each other. This prevents problems running the files or applications if a critical component is not installed. If you need to split an application’s files into multiple components, make sure that you make each component’s requirements the same, so that they all are either installed or not installed.

Follow the steps under [Section 5.4.5, “Configuring the Software Package Components,”](#) on page 256 and note the information you need to know for configuring the package components.

5.3.3 What Minimum Requirements Are Needed?

Minimum requirements establish whether a software package can be installed on the target machine. If these requirements are all met, you can install the software package on that server.

However, you can establish requirements for the software package as a whole, as well as for each package component. Therefore, if the package's requirements were all met, but some component requirements were not met, only part of the package would be installed.

Follow the steps under [Section 5.4.3, "Configuring the Server Software Package," on page 255](#) and note the information you need to know for configuring the software package.

5.4 Setting Up Server Software Packages

To set up a software package for distribution, perform the following tasks in order:

1. ["Setting Up Multiple-Workstation Management for Server Software Packages" on page 251](#)
2. ["Creating a Server Software Package" on page 255](#)
3. ["Configuring the Server Software Package" on page 255](#)
4. ["Creating the Software Package Components" on page 256](#)
5. ["Configuring the Software Package Components" on page 256](#)
6. ["Compiling a Software Package" on page 267](#)
7. ["Distributing the Software Package" on page 267](#)

5.4.1 Setting Up Multiple-Workstation Management for Server Software Packages

If you want to manage your software packages from multiple workstations, do the following in order to set up managing the replication of a master copy of the `snapinprefs.ser` file to multiple workstations; otherwise, continue with [Section 5.4.2, "Creating a Server Software Package," on page 255](#).

1. ["Setting Up the Master Snapinprefs.ser File" on page 251](#)
2. ["Creating and Using the ConsoleOne Batch File" on page 252](#)

Setting Up the Master Snapinprefs.ser File

For the following instructions, select any workstation that you use for managing software packages. If you have already created software packages using a workstation, select that workstation so you do not lose any software package information stored in the workstation's `snapinprefs.ser` file.

- 1** Map a drive to the server where you want to store your `.spk` and related `.cpk` files.

This drive letter should be one that can be used by all of the other workstations you use to manage software packages. This drive letter is written to the `snapinprefs.ser` file as part of the path information for each listed `.spk` file, so it should be a fixed drive letter that all workstations use.

The drive letter is also used in the batch file that you use to start ConsoleOne, which provides each workstation access to the same `.spk` file locations.

- 2** If you already have Server Software Package objects created by this workstation, skip to [Step 5](#).

or

If you have not yet created any Server Software Package objects using this workstation, start ConsoleOne.

This version of ConsoleOne must have the Policy and Distribution Services snap-ins installed.

3 In the Server Software Package namespace, create a Server Software Package object.

You do not need to fully configure the Server Software Package object at this time. Just give the package a name and provide a location and filename for the .spk file. Make sure you use the drive mapping you used in [Step 1](#).

For information on creating software packages, see [Section 5.4, “Setting Up Server Software Packages,” on page 251](#).

4 Exit ConsoleOne.

This step is important to make sure that the `snapinprefs.ser` file is created for this workstation.

5 On the network server you use to store the master copy of the `snapinprefs.ser` file, create a directory named `\C1` at the root of the drive.

You can select any safe location on the server for the master `snapinprefs.ser` file.

The [batch file sample](#) provided below uses a directory named `\C1`. You can modify the batch file if you want to use a different directory name, and you can include path information; however, do not use variables for the root location.

For example,

```
\zenworks\clssp
```

could be used to replace the `\C1` directory name.

6 Copy the workstation’s `snapinprefs.ser` file from:

```
c:\documents and settings\user_ID\.consoleone (Windows 2000)
```

to the `\C1` directory on the network server.

This becomes the master `snapinprefs.ser` file that is updated with new .spk paths, provided you are using the batch file documented in [“Creating and Using the ConsoleOne Batch File” on page 252](#).

7 Continue with [“Creating and Using the ConsoleOne Batch File” on page 252](#).

Creating and Using the ConsoleOne Batch File

Review the following sections to understand, create, and use the batch file:

- ◆ [“Sample Batch File” on page 252](#)
- ◆ [“What the Batch File Does” on page 253](#)
- ◆ [“Creating Your Batch File” on page 253](#)
- ◆ [“Optional Modifications to the Batch File” on page 254](#)
- ◆ [“Using the ConsoleOne Batch File” on page 254](#)

Sample Batch File

```
@echo off
REM map a network drive
net use m: \\server1.servers.novell.com\vol1

REM create a backup copy of the workstation’s .ser file
```

```

copy "%USERPROFILE%\consoleone\snapinprefs.ser"
"%USERPROFILE%\consoleone\snapinprefs.tmp"

REM copy the master .ser to the workstation
copy m:\c1\snapinprefs.ser "%USERPROFILE%\consoleone\snapinprefs.ser"

REM start ConsoleOne
c:\Novell\ConsoleOne\1.2\bin\ConsoleOne.exe

REM batch file control returns after exiting ConsoleOne
REM copy the updated .ser to server
copy "%USERPROFILE%\consoleone\snapinprefs.ser" m:\C1\snapinprefs.ser

REM restore the backup copy of the workstation's .ser file
copy "%USERPROFILE%\consoleone\snapinprefs.tmp"
"%USERPROFILE%\consoleone\snapinprefs.ser"

REM delete the mapped network drive
net use m: /delete
@echo on

```

What the Batch File Does

- ◆ It maps a network drive for accessing the server where you are storing .spk and .cpk files.
- ◆ It uses the %USERPROFILE% Windows variable to locate the Server Management \consoleone directory. This variable is also used by Server Management to determine where it creates the \consoleone directory and writes the snapinprefs.ser file.
- ◆ It creates a backup .tmp copy of the snapinprefs.ser file.
- ◆ It copies the master snapinprefs.ser file from the \C1 directory on the server to the workstation's \consoleone directory.
- ◆ It starts ConsoleOne.
- ◆ After you have exited ConsoleOne, the batch file copies the updated snapinprefs.ser file from the workstation's \consoleone directory to replace the version in the \C1 directory on the server. This becomes the new master snapinprefs.ser file.
- ◆ It restores the backed-up copy of the snapinprefs.ser file from the .tmp file.
- ◆ It unmaps the drive letter to the server.

Creating Your Batch File

- 1** Copy the text from the above sample batch file into a text editor.
- 2** Replace the m: drive letter with one that each of your workstations has free. Make sure you do this wherever m: exists in the batch file.
- 3** Edit the net use m: \\server1.servers.novell.com\vol1 line by replacing it with the path to the server volume or shared folder of the server where you are storing the .spk and .cpk files.
- 4** Save the batch file on your workstation and give it a name, such as:
C1SSP.BAT
- 5** Copy this batch file to each workstation that you use to manage software packages.

Optional Modifications to the Batch File

- ◆ If you installed ConsoleOne to a different location on the workstation than the one indicated in the batch file sample, modify the `c:\novell\consoleone\1.2\bin\consoleone.exe` line to reflect the location of the `consoleone.exe` file on the workstation.

You should make this modification in each individual batch file copy on a workstation where the default ConsoleOne path was not used.

- ◆ This batch file can also be used by a workstation to start an instance of ConsoleOne that is installed on a server. Modify the `c:\novell\consoleone\1.2\bin\consoleone.exe` line to reflect the location of the `consoleone.exe` file on the server. Make sure the drive letter is the one being used for accessing the server (see [Step 1 on page 251](#)).

- ◆ If the `\consoleone` directory path is different between workstations because the `%USERPROFILE%` variable was not used, you need to edit any lines containing the variable, as necessary. Open the copy of the batch file on a workstation where the `%USERPROFILE%` variable was not used and edit the lines containing the variable to reflect the correct path to the `\consoleone` directory.

- ◆ If you created a directory other than `\C1` on the server, replace `\C1` wherever it exists in the batch file with the directory that you specified in [Step 5 on page 252](#).

- ◆ The batch file creates a `.tmp` version of the `snapinprefs.ser` file. This allows you to maintain the version of the `.ser` file on the workstation that existed before you used the batch file. However, if you want the workstation's version to always match the master version it copied to the server, remove the following two lines from the batch file:

```
copy "%USERPROFILE%\consoleone\snapinprefs.ser"  
"%USERPROFILE%\consoleone\snapinprefs.tmp"
```

```
copy "%USERPROFILE%\consoleone\snapinprefs.tmp"  
"%USERPROFILE%\consoleone\snapinprefs.ser"
```

- ◆ If you cannot use the same drive letter for all workstations, you can use the `%1` argument in the batch file, but only if you are using UNC paths for all of your `.spk` files. To do this, replace all occurrences of `m:` with `%1`. Then, when you execute the batch file from a command line, add the drive letter after the batch file's name. For example,

```
C:\SSP R:
```

causes the batch file to use R: as the drive for locating the master copy of the `snapinprefs.ser` file.

Using the ConsoleOne Batch File

- ◆ Before running this batch file, place a `snapinprefs.ser` file in the `\consoleone` directory of each workstation you use to manage software packages. The batch file assumes that the `.ser` file exists for copying and replacing.
- ◆ Before running this batch file, place your master copy of the `snapinprefs.ser` file in the `\C1` directory of the server where you have stored the software package files. The batch file assumes that this `.ser` file exists for copying and replacing.
- ◆ Run this batch file any time you plan to add, delete, or compile software packages.

- ◆ You do not need to use the batch file when you view or edit the properties of software packages. The add, delete, and compile functions are the only actions that causes the `snapinprefs.ser` file to be updated.

Continue with [Section 5.4.2, “Creating a Server Software Package,”](#) on page 255.

5.4.2 Creating a Server Software Package

- 1 In ConsoleOne, right-click the Server Software Packages namespace, then click *New Package*.

The Create New Server Software Package Wizard opens.

- 2 Read the information on the first dialog box, then click *Next*.

- 3 Provide a name for the software package.

Make this a descriptive name. It is displayed in ConsoleOne under the Server Software Packages object.

IMPORTANT: Do not use double-byte characters in the software package name. This causes an error in any report you run on the software package.

- 4 Because software packages are file-based, provide the full path and filename, including the `.spk` extension.

If you don't enter the extension, you are prompted to add it.

You can also use UNC paths.

You can store the `.spk` files on a workstation or server. The `.spk` files is typically below 100 KB in size. However, compiled software packages (`.cpk` files) can be in the hundreds of megabytes. For information on storing `.spk` and `.cpk` files, see [Section 5.2.9, “Managing Server Software Packages,”](#) on page 243.

WARNING: Software package full paths and filenames are drive-dependent. A software package can become unusable if you change drive mappings after creating the package. Make sure your entry in this field is not changed. However, if you used a UNC path, this is not an issue.

- 5 Click *Finish*.

- 6 Continue with [Section 5.4.3, “Configuring the Server Software Package,”](#) on page 255.

5.4.3 Configuring the Server Software Package

After a software package has been created, you need to configure it by setting the prerequisites for installation of the files and applications contained in the package.

To configure a package:

- 1 In ConsoleOne, right-click a software package, then click *Properties*.

The *Identification* tab should be displayed. If not, select it.

The *Name* field should display the name you gave the package when you created it.

- 2 Provide a useful description for the software package.

- 3 If you don't want to be able to roll back to the older version of the server file or application after installing the newer version, click *Disable Rollback*. However, this is not recommended.
For information on rolling back software package installations, see [Section 5.2.11, "Rolling Back Software Package Installations,"](#) on page 249.
- 4 Select the *Requirements* tab.
- 5 Click *Add*, then select a requirement:
 - [Section E.1, "Operating System,"](#) on page 413
 - [Section E.2, "Memory \(RAM\),"](#) on page 416
 - [Section E.3, "Disk Space,"](#) on page 416
 - [Section E.4, "SET Commands,"](#) on page 417
 - [Section E.5, "Registry,"](#) on page 418
 - [Section E.6, "File,"](#) on page 418
 - [Section E.7, "Products.dat,"](#) on page 418
- 6 Repeat [Step 5](#) for each requirement.
- 7 If you want to use variables to customize the installation, select the *Variables* tab, then click *Add*.
- 8 Provide the variable name and value.
For information on variables, see [Section 9.6, "Using Variables to Control File Extraction,"](#) on page 349.
- 9 Repeat [Step 7](#) and [Step 8](#) for each variable.
- 10 Click *OK* when you have finished configuring.
If you click *Cancel*, none of the configuration changes on any of the tabs are saved.
- 11 Continue with [Section 5.4.4, "Creating the Software Package Components,"](#) on page 256.

5.4.4 Creating the Software Package Components

After you have created and configured a software package, you need to create the components of the package, including the individual files or applications for the package.

To create the software package components:

- 1 In ConsoleOne, right-click a software package (in the left pane), then select *New Component*.
- 2 Provide the name of the component as you want it to be displayed in ConsoleOne, then click *OK*.
The component is displayed as named under the Software Package object.
- 3 Repeat these steps for each component needed.
- 4 Continue with [Section 5.4.5, "Configuring the Software Package Components,"](#) on page 256.

5.4.5 Configuring the Software Package Components

After you have created the software package components, you need to configure the prerequisites for each, including identifying the files or applications for the component.

Package components can each have the same prerequisites, which can save time and minimize user error.

To configure a component:

- 1 In ConsoleOne, right-click a component, then click *Properties*.

The *Identification* tab should be displayed. If not, select it.

- 2 Provide a useful description for the component.
- 3 Select a further action for the software package to perform after the installation process has finished from the *After package installation is complete* drop-down list.
- 4 To continue configuring the component, see each of the following that you might need to configure:

“Requirements” on page 257

“Pre-Installation Load/Unload” on page 258

“Pre-Installation Script” on page 258

“Local File Copy” on page 259

“Copy File” on page 260

“Text Files” on page 263

“SET Commands” on page 264

“Registry Settings” on page 265

“Products.dat” on page 265

“Post-Installation Unload/Load” on page 266

“Post-Installation Script” on page 266

Do not click *OK* on this component’s property page until you have finished configuring all of the above items, as needed.

- 5 Click *OK*.

If you click *Cancel*, none of the configuration changes on any of the tabs are saved.

- 6 Continue with [Section 5.4.6, “Compiling a Software Package,” on page 267](#) to ready your software package for distribution.

Requirements

To specify requirements for installing the server files or applications:

- 1 While displaying the properties of the software package component, select the *Requirements* tab, then click *Add*.
- 2 Select any of the following requirement items:

[Section E.1, “Operating System,” on page 413](#)

[Section E.2, “Memory \(RAM\),” on page 416](#)

[Section E.3, “Disk Space,” on page 416](#)

[Section E.4, “SET Commands,” on page 417](#)

[Section E.5, “Registry,” on page 418](#)

[Section E.6, “File,” on page 418](#)

[Section E.7, “Products.dat,” on page 418](#)

For further instructions on configuring an item, see one of the above items.

Continue with the next item to configure before clicking *OK*.

Pre-Installation Load/Unload

To configure certain NLM files or processes to load or unload before installing the software package on a server:

- 1 While displaying the properties of the software package component, select the *Pre-Installation* tab, then click *Load/Unload*.
- 2 Click *Add*.
- 3 Select one of the following:

[Section D.1, “Load NLM/Process,” on page 411](#)

[Section D.2, “Load Java Class,” on page 411](#)

[Section D.3, “Unload Process,” on page 412](#)

[Section D.4, “Start Service,” on page 412](#)

[Section D.5, “Stop Service,” on page 412](#)

For further instructions on configuring an item, see one of the above items.

IMPORTANT: If you select a process to be loaded by the software package, and it is already running on the target server, the package installation fails and is rolled back (if rollback is enabled). If the process requires intervention to unload, you must remember to unload it manually before installing the software package.

To make sure that a process is not already loaded when you are including it in the software package, add an unload option for that process before adding the load option—but only if the process does not require user input from the keyboard to unload it.

- 4 Repeat **Step 1** through **Step 3** for each NLM or process to be included.
- 5 Use the arrow keys to arrange the order to execute the NLM files and the processes.

Continue with the next item to configure before clicking *OK*.

Pre-Installation Script

To configure running server scripts before installing the software package on a server:

- 1 While displaying the properties of the software package component, select the *Pre-Installation* tab, then click *Script*.
- 2 Click *Add*.
- 3 Provide the script name.
- 4 Select the script type (NCF, NetBasic, PERL).

IMPORTANT: NetBasic is not supported on NetWare 6.5 servers.

- 5 Enter the script text.

WARNING: If a software package passes all requirements and executes the script, processing done by the script cannot be undone by rollback.

- 6 Repeat **Step 2** through **Step 5** for each script to be added.
- 7 Use the arrow keys to arrange the order to execute the scripts.

Continue with the next item to configure before clicking *OK*.

Local File Copy

The Local File Copy component enables copying of files on a server from one location to another using a software package. You can either copy or move the files.

To configure the Local File Copy component:

1 While displaying the properties of the software package component, select the *Local File Copy* tab.

2 Click *Add*.

Local File Copy #1 is the default name. You can edit that name.

3 Fill in the fields:

Source path: Provide the full path where the files to be copied are located.

You can use wildcards in the path:

* = any number of characters

? = any single character in that position

??? = any characters in those positions

Target path: Provide the full path where the copied files are to be placed.

You can use wildcards in this path. This path does not need to mirror the source path. However, you could mirror an existing target path.

Include subdirectories: Includes all subdirectories and their files beginning from the directory at the end of the path; otherwise, only the files in the directory at the end of the path are copied.

Maintain attributes: Maintains the file attributes in the target's file system that exist in the source's file system.

Overwrite Destination Files: Overwrites files of the same name in the destination directories, regardless of differences in file dates. If you do not select this option, files of the same name are not replaced.

Maintain trustees: Maintains the file's trustee attributes.

When a file is locked: Select one or both:

- ♦ **Retry __ times:** Retries overwriting a locked file the number of times you select before failing to replace the file. Leave this check box deselected to not replace locked files on the target file system.
- ♦ **Kill connection of open files:** (NetWare only) Attempts to kill the connection of locked files so that they can be overwritten. This applies only to files being extracted, not to files being accessed to build the Distribution. If a file belonging to a Distribution is locked when the Distribution is being built, the build fails. Server and NLM connections cannot be killed.

Error processing: *Fail On Error* is selected by default. This stops the file copying process when an error is encountered in copying. To continue file copying when an error is encountered, select *Continue On Error*.

Operation: Sets whether to copy or move the files identified in the *Source Path* field.

Continue with the next item to configure before clicking *OK*.

Copy File

You can configure the Copy File component to control how files are copied during installation of a software package. This includes adding files to existing directories, creating new directories, adding files and subdirectories to the new directories, and deleting existing files and directories.

A file group is a root item for the component's expandable tree structure. You can have multiple file groups for the Copy File component. A file group is a set of related directories and files. File groups are top-level items and cannot contain other file groups.

The other structure items are directory and file, which are contained within a file group. Directories can contain other directories or files, but not file groups.

IMPORTANT: When you add a file group or directory, you are creating the target paths where the files are to be copied, not the source paths of the files. The source paths are automatically accounted for as you select your source files or directories.

To configure copying files during installation of the software package:

- 1** While displaying the properties of the software package component, select the *Copy File* tab.
- 2** To create your first file group, do the following:
 - 2a** Click the down-arrow on the drop-down list next to the *Add* button, select *Add File Group*, then click *Add*.

Because files and directories must be contained within file groups, you are prompted to create a file group the first time you click *Add*, regardless of the type you are attempting to add.

You should create one file group for each specific target location. For example, `c:\files`, `c:\data\accounting`, and `c:\data\personnel` could be different locations on a C: drive where you want to copy different groups of unrelated files.
 - 2b** Name the file group, then provide its target path.

The file group's target path specifies the base path from where all directories and files within the group are installed.
 - 2c** To specify what to do when a file group location is locked, select the check box for one of the following:
 1. Retry (enter the number of retry times)
 2. Kill Connection of Open Files (NetWare only)
 3. Fail With Error

Retries are about 5 seconds apart. Therefore, 12 retries would take about one minute.
- 3** To create a target directory under a file group or another directory, select the file group or directory, in the drop-down box select *Add Directory*, click *Add*, then do the following:
 - 3a** Because *Directory* is the default directory name, to rename the directory, right-click *Directory*, click *Rename*, type the desired directory name, then press Enter.

When entering information into this field, you must press Enter for the change to be saved.

To match an existing target directory for deleting or copying files, you must enter the exact name.

IMPORTANT: If you provide an existing directory name and that directory is marked as Read Only on the destination server's file system, the Software Package Distribution fails

when the Subscriber tries to extract the Distribution, because it cannot write to that directory. Therefore, you must know the attributes of existing target directories and remove their Read Only directory attributes.

You must create the same directory structure in the File Copy component as exists in the target location so that the directory name you provide here is in the same sequence in the path.

- 3b** To determine whether to create or delete the directory, select the mode from the *Copy Mode* drop-down list.

Create: If you select *Create* and the directory does not exist, the directory is created. If you select *Create* and the directory does exist, the directory is not created, and no error is encountered.

Delete: If you select *Delete* and the directory exists, the directory is deleted, including any subdirectories and files under it. If you select *Delete* and the directory does not exist, the directory is not deleted, and no error is encountered.

WARNING: If you plan to set the *Copy Mode* as Delete for any directories you add, and you do not want any parent directories that you have added to also be deleted, place those parent directories in the *Target Path* field of the file group. For example, if you want to delete `c:\winnt\cookies`, but do not want to delete the `\winnt` directory, enter `c:\winnt` in the *Target Path* field, click *Add* to enter the `\cookies` directory in the tree structure, then click *Delete* for the *Copy Mode* field. For example:

Target = `c:\winnt`

Tree structure = `cookies`

causes only `cookies` and all of its files and subdirectories to be deleted.

Conversely, both the `\winnt` and `\cookies` directories are deleted if you enter `c:\` in the *Target Path* field, click *Add* to enter the `\winnt` directory in the tree structure, click *Add* to enter the `\cookies` directory under `\winnt` in the tree structure, then click *Delete* for the *Copy Mode* field.

For example:

Target = `c:\`

Tree structure = `winnt\cookies`

causes `winnt` and all of its files and subdirectories to be deleted.

- 4** To add files or source directories under a file group or directory in the tree structure, select a file group or directory, in the drop-down box select *Add File*, click *Add*, then do the following:

- 4a** Select the files or directories using the Open dialog box.

These directories and files are displayed directly under the file group or directory you selected in [Step 3](#).

For the destination server's file system, attributes of the copied files and directories are not maintained. For more information, see [Step 4c](#).

If you selected a directory on the Open dialog box, it is not displayed expanded. Click the plus signs to expand the existing structure under the directory that you added.

In the Open dialog box, the Recurse Directories option is selected by default. To only select files in this directory, select the Recurse Directories check box to disable it and none of the subdirectories are selected.

To exclude files or subdirectories from being selected, select the Exclude Selected Subdirectory option, select the files or directories to be excluded (use Shift and Ctrl for multiple select), then click Open.

If you exclude files or subdirectories, it does not remove them from the file system. It only prevents them from being selected.

For information on removing files or subdirectories from the tree structure after adding files and directories, see [Step 8](#).

4b To configure a subdirectory that was added, do the following:

- ◆ Select the subdirectory, then select the *Copy Mode* (whether to Create or Delete the directory).

WARNING: When you set the *Copy Mode* to Delete, it causes deletion of the target directory and all of its files and subdirectories.

- ◆ To rename a subdirectory that was added, right-click the subdirectory, click *Rename*, type a new directory name, then press Enter.

When entering information into this field, you must press Enter for the change to be saved.

If you rename a directory that was selected through the Open dialog box, make sure that the new name meets your expectations for the target location.

Because only selected files have their path remembered for copying, renaming a directory does not affect file selection. In other words, you can give a target directory a different name than its source, and still have the same files copied under it.

4c To configure an added file, select the file, then do the following:

- ◆ To determine the file's copy mode, select a mode from the *Copy Mode* drop-down list.

You must select an option for every file. You can select multiple files where you want the mode to be the same.

The options are: Copy Always, Copy If Exists, Copy If Does Not Exist, Copy If Newer, Copy If Newer and Exists, and Delete.

WARNING: When you set the *Copy Mode* to Delete, it causes deletion of the selected file from the target server.

- ◆ Select the check box for each attribute that should apply to the selected files.

Attributes do not default. You must set them for the destination server. They are not carried over from where you obtained the file.

IMPORTANT: Do not select all of the attributes for a file, or an exception is thrown on the server.

When setting the attribute of an executable file, set it to Read Only. Do not set it to Execute. If you mark a file as Execute, the NetWare[®] CLIB API does not allow you to change it to a different attribute. To change the attribute from Execute to Read Only after the software package has been installed, you need to manually delete the file, replace it, then set its attribute again.

5 To create another file group, do the following:

5a Click the down-arrow on the drop-down list next to the *Add* button, select *Add File Group*, then click *Add*.

It doesn't matter what you have selected in the tree structure; the file group is automatically placed at the first tree level, equal to any other file groups that are displayed.

5b Name the group.

5c Provide its target base path.

5d To indicate what to do when a group location is locked, select the check box for one of the following:

1. Retry (enter the number of retry times)
2. Kill Connection of Open Files (NetWare only)
3. Fail With Error

6 Repeat **Step 5** or each additional file, directory, or file group to be added.

7 If you want the file groups to be copied in a particular order, use the arrow keys to arrange the order of the file groups.

The arrows are dimmed if the file group you have selected has no valid up or down movement available to it.

8 To remove a file group, directory, or file, select it, then click *Remove*.

You can use the *Remove* button to prune the tree structure of unwanted files or directories.

You can use the Shift and Ctrl keys to select multiple items for removal.

IMPORTANT: If you remove a file group or directory, all files and directories displayed below it are also removed, but only from this tree structure, not from the source file system.

Continue with the next item to configure before clicking *OK*.

Text Files

To configure making changes to text files during installation of the software package:

1 While displaying the properties of the software package component, select the *Text Files* tab.

2 Click *Add*.

After one text file has been added, you are given the opportunity to select whether you are adding another text file or adding another change item for the selected text file.

To add another text file: Select *Text File*. It does not matter which text file or change item is selected in the left pane—the text file is added to the far left level.

To add another change to a text file: In the left pane select the text file for the change, click *Add*, then select *Change*. The change item is added under the selected text file.

3 If you are adding a text file, provide the name of the text file.

4 Accept the default name (such as Change #1) or rename it.

If you are adding a text file, click *OK*.

5 Click the down-arrow for the *Change Mode* field, then select the change mode from the drop-down list.

6 Click the down-arrow for the *Search Type* field, then select the search type from the drop-down list.

7 Enter the exact search string.

- 8 Select the check box if you want the string search to be case sensitive.
- 9 To find all occurrences of the search string, select the check box (default); otherwise, deselect the check box to find only the first occurrence.
- 10 Click the down-arrow for the *Result Action* field, then from the drop-down list, select the action that should result if a string is matched.
- 11 If you are replacing a string or entering a new one, enter the text in the *New String* text box.
- 12 Repeat **Step 2** through **Step 11** for each text file to add or each change to be made.
- 13 To reorder the text files and change items, use the arrow keys.

Continue with the next item to configure before clicking *OK*.

SET Commands

For NetWare only.

To configure the target server's SET commands:

- 1 While displaying the properties of the software package component, select the *SET Commands* tab.
- 2 Click *Add* to open the NetWare Server SET Commands Wizard.
- 3 Select the server containing the SET commands, then click *Next*.

IMPORTANT: The Server Management and Java must be running on the server where you want to obtain the SET commands.

- 4 Select all of the SET commands you want to configure for the target server.
You can select whole categories by selecting the check box for the category, or click the plus sign to expand a SET command category and select the check boxes for individual SET commands to be included.

WARNING: Do not select the Set Developer Option SET command and change Off to On. This parameter is meant to help developers debug server abends. It disables some operating system checking to prevent certain abends from occurring. Also, if the Set Developer Option is enabled, running NCP™ scripts that require keyboard entry could abend the server.

- 5 Click *Finish* when you have completed selecting SET commands.
The selected SET commands are now displayed in the *SET Commands* tab for the file or application component.
- 6 To edit a SET command, click its plus sign to expand its attributes.
- 7 To edit an attribute, select the attribute, then click *Edit*.
A dialog box is displayed where you can make changes to the attribute.
- 8 Repeat **Step 7** for each attribute to edit for a given SET command.
- 9 Repeat **Step 6** through **Step 8** to edit another SET command's attributes.

Continue with the next item to configure before clicking *OK*.

Registry Settings

To configure registry changes for either NetWare or Windows servers:

- 1 While displaying the properties of the software package component, select the *Registry Settings* tab, then click HKEY_LOCAL_MACHINE.

HKEY_LOCAL_MACHINE is a Windows registry key. For NetWare, HKEY_LOCAL_MACHINE is also recognized by Server Management as the equivalent to My Server. Therefore, you can use this key for editing both NetWare and Windows registries.

- 2 Click *Add*.
- 3 Select from the following:

Section F.1, “Key,” on page 421

Section F.2, “Binary,” on page 422

Section F.3, “Expand String,” on page 422

Section F.4, “(Default),” on page 422

Section F.5, “DWord,” on page 423

Section F.6, “Multi-Value String,” on page 423

Section F.7, “String,” on page 423

For further instructions on configuring an item, see one of the above items.

- 4 Repeat **Step 2** and **Step 3** for each registry entry to be made.
- 5 Use the arrow keys to arrange the order in making registry entries.

Continue with the next item to configure before clicking *OK*.

Products.dat

For NetWare only.

The `products.dat` file can be updated by your software package so that future updates can identify the most recently installed version of the file or application.

WARNING: Modifying `products.dat` could prevent something from running or being installed on the NetWare server. Never modify any entries supplied by Novell.

To determine which action to take for `products.dat`:

- 1 While displaying the properties of the software package component, select the *Products.dat* tab.
- 2 Select one of the following:

| Option | Description |
|------------------------|---|
| Add | Adds a new entry |
| Modify Existing Entry | Searches for a matching ID and modifies the version and description |
| Replace Existing Entry | Searches for a specific ID and replaces it with a new one |
| No Action | This is the default. Nothing is done to <code>products.dat</code> |

- 3** If you selected *Add*:
 - 3a** Provide the ID of the item to add.

This is case sensitive. The item is the ID of the new product for the `.dat` file.
 - 3b** Provide the exact version number to add.
 - 3c** Provide the description to add.
- 4** If you selected *Modify Existing Entry*:
 - 4a** Provide the ID of the item to search for (case sensitive).
 - 4b** Provide the new version number.
 - 4c** Provide the new description.
- 5** If you selected *Replace Existing Entry*:
 - 5a** Provide the ID of the item to search for (case sensitive).
 - 5b** Provide the exact version number to match.
 - 5c** Provide the new ID.
 - 5d** Provide the new version.
 - 5e** Provide the new description.

Continue with the next item to configure before clicking *OK*.

Post-Installation Unload/Load

To configure certain NLM files and processes to load or unload after installing the software package on a server:

- 1** While displaying the properties of the software package component, select the *Post-Installation* tab, then click *Load/Unload*.
- 2** Click *Add*.
- 3** Select one of the following:
 - Section D.1, “Load NLM/Process,” on page 411
 - Section D.2, “Load Java Class,” on page 411
 - Section D.3, “Unload Process,” on page 412
 - Section D.4, “Start Service,” on page 412
 - Section D.5, “Stop Service,” on page 412Select an item for further instructions on configuring it.
- 4** Repeat **Step 2** and **Step 3** for each NLM or process to be included.

Continue with the next item to configure before clicking *OK*.

Post-Installation Script

To configure running NetWare server scripts after installing the software package on a server:

- 1** While displaying the properties of the software package component, select the *Post-Installation* tab, then click *Script*.
- 2** Click *Add*.
- 3** Provide the script name.

- 4 Select the script type (NCF, NetBasic, PERL).

IMPORTANT: NetBasic is not supported on NetWare 6.5 servers.

- 5 Enter the script text.

WARNING: If a software package passes all requirements and executes the script, processing done by the script cannot be undone by rollback.

- 6 Repeat **Step 2** through **Step 5** for each script to be added.
- 7 Use the arrow keys to arrange the order to execute the scripts.

Continue with the next item to configure before clicking *OK*.

5.4.6 Compiling a Software Package

Your software packages (.spk files) cannot be installed by Policy and Distribution Services until they have been compiled and have the .cpk extension.

To compile a software package:

- 1 In ConsoleOne, right-click a software package, then click *Compile Package*.

The Compile Server Software Package Wizard opens.

- 2 Read the information on the first dialog box, then click *Next*.
- 3 Provide a name and path for the compiled software package (using the .cpk extension), then click *Next*.

Select a location where free disk space is adequate for the .cpk file. Compiled software packages (.cpk files) are generally much larger than the uncompiled (.spk) counterparts.

IMPORTANT: If you provide the path and filename of the .spk when prompted for the compiled (.cpk) filename, the .spk is overwritten and can no longer be edited. Therefore, be sure to use the .cpk extension when naming the compiled version.

The compiling process could take some time, depending on how many files are involved.

- 4 When compiling has completed, click *Finish*.
- 5 Continue with **Section 5.4.7, “Distributing the Software Package,”** on page 267 to distribute your software package (.cpk).

5.4.7 Distributing the Software Package

After a software package is ready for distribution, you can distribute it in the following ways:

- ♦ Use Tiered Electronic Distribution (see **Chapter 3, “Tiered Electronic Distribution,”** on page 85 for instructions on distributing through Tiered Electronic Distribution)
- ♦ Manually copy the software package file (.cpk) to the server and run it from the server’s console prompt using the PACKAGE command (see **Appendix C, “Server Console Commands,”** on page 405 for instructions on using the command)

After a software package is installed on a target server, you might need to reboot the server. For example, if tcpip.nlm is modified by the package, it cannot be downed—you must reboot the

server to run that NLM again. However, you could have the software package cause the server to come down and restart automatically.

5.5 Using Server Software Packages to Delete Directories on Servers

If you want to delete certain directories from a number of different network servers (NetWare, Windows, Linux, and Solaris), you normally do not have an automated method for performing this task. However, if you are using ZENworks 7 Server Management Policy and Distribution Services, the Server Software Packages feature of Server Management provides the capability for you to delete specified directories from any Subscriber server's file system.

To automate the deletion of specified directories on multiple servers, you first set up path variables (if necessary), create a Server Software Package in its namespace in ConsoleOne, compile the software package, then distribute the package using Tiered Electronic Distribution. No further user intervention is required.

Do the following in order to create a software package that deletes specified directories on a server:

1. [“Setting Up Variables for Use With the Server Software Package” on page 268](#)
2. [“Creating the Server Software Package” on page 269](#)
3. [“Creating and Configuring the Server Software Package Component” on page 269](#)
4. [“Compiling the Server Software Package” on page 271](#)
5. [“Manually Testing that the Directories Have Been Deleted” on page 271](#)
6. [“What’s Next” on page 271](#)

5.5.1 Setting Up Variables for Use With the Server Software Package

Before you create the software package, you must set up the variables in your Subscriber objects' properties if you are using variables in paths (for instance, if your target servers have different operating systems, like NetWare and Windows).

1 Identify the directories to be deleted:

- 1a** Identify the root of the path, such as its volume name (NetWare), drive letter (Windows), or /var/opt/novell/zenworks (for Linux and Solaris). For example, `data:`.
- 1b** Identify the rest of the path, including the parent directory to the directories to be deleted, such as `\zenworks\pds\ted\dist` where `dist` is the parent directory.
- 1c** Identify the directories to be deleted, such as `olddist.Distributions.ZENworks.Novell`.

The resulting full path and directory to be deleted would be:

```
data:\zenworks\pds\ted\dist\olddist.Distributions.ZENworks.Nove  
ll
```

You might have varying path elements from server to server. You should use variables as necessary to allow for those differences (see [Step 2](#) and [Step 3](#)).

- 2** In ConsoleOne, create a variable to represent `data:`, `D:`, or `/var/opt/novell/zenworks` for each Subscriber where the directories to be deleted reside, such as `DELETEDDIRROOT`.

If you name a directory to be deleted that does not exist on a target server, nothing is done for that directory on that server.

You can also define variables globally using the Tiered Electronic Distribution policy. There, you should define the default value for a variable and allow the exceptions to be defined in the applicable Subscriber objects' properties.

- 3 In ConsoleOne, create a Subscriber variable to represent where any path elements are different for the Subscriber server.

For more information on variables, see [Chapter 9, "Variables," on page 341](#).

- 4 Repeat [Step 2](#) and [Step 3](#) as necessary.
- 5 Continue with [Section 5.5.2, "Creating the Server Software Package," on page 269](#).

5.5.2 Creating the Server Software Package

- 1 In the left pane in ConsoleOne where the ZENworks 7 Server Management snap-ins have been installed, right-click the *Server Software Packages* namespace.
- 2 Click *File > New > Software Package* to start the Create New Server Software Package Wizard.
- 3 Click *Next*.
- 4 Provide a name for the software package, such as Delete Old Directories.
- 5 Specify a path and filename for the software package template file (.spk), such as `c:\temp\deletedirs.spk`.

IMPORTANT: If you save your .spk file to a network server, use a UNC path so that you still have access to that software package file if your drive letters change.

You can also save your .spk files to a workstation or server, because the .spk file sizes do not become large. For this particular type of software package (where you are only giving instructions for deleting directories and are not compiling data files), the .cpk (compiled software package) version is similar in size. Therefore, for management purposes, you might want to save these .spk files and their corresponding .cpk files in the same location, which can be on a workstation or server.

- 6 Click *Finish*.
- 7 If necessary, click the plus sign to expand the *Server Software Package* namespace to view the new package.
- 8 Continue with [Section 5.5.3, "Creating and Configuring the Server Software Package Component," on page 269](#).

Unless otherwise instructed, you should perform the steps in the subsequent sections from the same instance of ConsoleOne you used in the above steps, because your .spk files are accessible from there.

5.5.3 Creating and Configuring the Server Software Package Component

- 1 Right-click the software package object that you just created, then select *New Component*.
- 2 Provide a name for the component, such as Delete Directories.
- 3 If necessary, click the plus sign to expand the Server Software Package object.

- 4 Right-click the component and select *Properties*.
- 5 Select the *Copy File* tab.
- 6 Click the drop-down list button next to the *Add* button, then select *Add File Group*.
- 7 Click *Add*.
- 8 Provide a name for the file group, such as Delete Working Directories.
- 9 In the *Group Target Path* field, specify the name of the variable that you created containing the location of the directories to be deleted, and add any path information that is not contained in the variable; however, do not specify the name of the directory to be deleted as part of that path.
For example, if the location for the directories to be deleted is the same for all target servers, specify the actual volume (NetWare) or drive (Windows) with the path information (which can also contain variables).

However, if you need to use variables because the server operating systems are different, then specify the variable name (within the % symbols) plus the full path (which can also contain variables) to the directory just above the directories to be deleted. For example, %DELETEDDIRROOT% (variable name) and %TARGET%\pds\ted\dist (full path to the parent directory of the directories you want to delete).

IMPORTANT: When using variables, the path you provide must be the directory containing the directory to be deleted. In [Step 11](#) you add the actual directory names to be deleted.

- 10 Click *OK* to exit the dialog box.
- 11 Click the drop-down list button again and select *Add Directory*.
Make sure you first select the tree item under which you want to add this directory.
- 12 Click *Add*.
- 13 To change the name (“Directory”) that defaults in the tree structure to the actual directory name that you want deleted (such as `olddist.Distributions.ZENworks.Novell`), edit the directory name and press the Enter key to save the change.
If you do not press the Enter key, “Directory” is displayed again. The *Rename* button allows you to edit the directory name.
- 14 Click the drop-down list button next to the *Copy Mode* combo box, then select *Delete*.
- 15 Click *Apply*.
- 16 Repeat [Step 10](#) through [Step 15](#) for each directory you want this software package to delete using this component’s file group.
You can start at [Step 6](#) to add other file groups, or from [Step 1](#) to add a new component. You might want to repeat from these steps if you cannot add all of your directories to be deleted under the file group that you created in [Step 6](#).
- 17 When finished configuring the software package component, click *OK* or *Close*.

Using the examples from the above steps, you would have entered:

```
%DELETEDDIRROOT%
```

and

```
%TARGET%\pds\ted\dist
```

and

```
olddist.Distributions.ZENworks.Novell
```

in order to delete the directories having the following paths:

```
data:\zenworks\pds\ted\dist\olddist.Distributions.ZENworks.Novell
d:\zfs\zenworks\pds\ted\dist\olddist.Distributions.ZENworks.Novell
```

18 Continue with [Section 5.5.4, “Compiling the Server Software Package,”](#) on page 271.

5.5.4 Compiling the Server Software Package

You now have a `.spk` file that serves as the template for what you want to delete. You need to compile this `.spk` file into a `.cpk` file.

- 1 Right-click the software package, such as Delete Old Directories.
- 2 Select *Compile* to start the Compile Software Package Wizard.
- 3 Click *Next* on the first page of the wizard.
- 4 Provide the full path and filename for the `.cpk` file that you are generating.

IMPORTANT: Do not use the `.spk` extension for this filename, or your template file could be overwritten by its compiled version if they are stored in the same location. This would prevent you from making further edits to the software package. You can use the same filename, such as `DELETEDIRS`, but you should use only the `.cpk` filename extension.

- 5 Click *Next*, then click *Finish*.
- 6 Continue with [Section 5.5.5, “Manually Testing that the Directories Have Been Deleted,”](#) on page 271.

5.5.5 Manually Testing that the Directories Have Been Deleted

The software package is now ready for sending as a Software Package Distribution. However, for testing, you can manually process the software package on one of the target servers to determine that the directories were deleted as intended.

- 1 On a server where you want to delete a directory, create a directory that is contained in your software package (such as `olddist.Distributions.ZENworks.Novell`) under `\zenworks\pds\ted\dist`.
- 2 Copy the `.cpk` file (for example, `deletedirs.cpk`) to the `\temp` directory on that server.
- 3 At the server’s ZENworks Server Management console prompt, enter the `PACKAGE PROCESS` command to process the software package.

For example, if it is a NetWare server, at the ZENworks Server Management prompt you should enter:

```
package process data:\temp\deletedirs.cpk
```

Server Management processes the package and report that it has finished processing. Check the server’s file system to see that the `\olddist.Distributions.ZENworks.Novell` directory, or the directories you specified, were deleted.

5.5.6 What’s Next

After you are satisfied with the result of your test, you can distribute the `deletedirs.cpk` file using Tiered Electronic Distribution to all your target Subscriber servers with your new Software Package Distribution in order to delete directories on your Subscriber servers’ file systems.

Desktop Application Distribution

6

Novell® ZENworks® Server Management provides Policy and Distribution Services integration with ZENworks Desktop Management's Novell Application Management.

Desktop Application Distributions can be sent to only Linux, NetWare®, and Windows servers. This Distribution type is not supported on Solaris servers.

The following sections provide information on understanding, setting up, and using the integration between ZENworks Server Management and ZENworks Desktop Management:

- ◆ [Section 6.1, “Understanding Desktop Application Distributions,” on page 273](#)
- ◆ [Section 6.2, “Requirements,” on page 288](#)
- ◆ [Section 6.3, “Creating a Desktop Application Distribution,” on page 289](#)
- ◆ [Section 6.4, “Rebuilding Desktop Application Distributions,” on page 296](#)
- ◆ [Section 6.5, “Cleaning Up Desktop Application Distribution Files,” on page 297](#)
- ◆ [Section 6.6, “Sending Desktop Application Distributions Tree-To-Tree,” on page 298](#)

6.1 Understanding Desktop Application Distributions

Server Management allows you to solve geographic, workload, and redundancy issues for applications distributed by Novell Application Launcher™ that might otherwise require much of your time in manual configuration work in Desktop Management. Review the following sections to see how Server Management can help you to automate much of your desktop application work.

- ◆ [Section 6.1.1, “The Purpose of Desktop Application Distributions,” on page 273](#)
- ◆ [Section 6.1.2, “Distributed Application Issues,” on page 275](#)
- ◆ [Section 6.1.3, “Miscellaneous Issues,” on page 286](#)

6.1.1 The Purpose of Desktop Application Distributions

- ◆ [“Applications in Desktop Management” on page 273](#)
- ◆ [“Distributed Applications in Server Management” on page 274](#)

Applications in Desktop Management

In Desktop Management, you can create Application objects so that users or workstations can receive their applications through Novell Application Launcher. An Application object contains pointers to the files belonging to the application, and also contains configuration parameters for how the application is to be installed and configured on the desktop.

In Desktop Management, the files belonging to an application can exist on any server, and the related Application object can exist anywhere in the tree. Therefore, for a workstation to receive an application through Novell Application Launcher, the application's files are copied from a server and installed on the workstation.

However, problems can arise for the Desktop Management administrator, such as:

- ♦ **Network traffic:** Many users or workstations can create heavy network traffic (especially across slower WAN links) to obtain their applications

To address the geographic issue of heavy network traffic, if you use only Desktop Management, you would need to do a lot of manual work. You would have to re-create and custom-configure the Application objects multiple times and copy their files to the various servers that would locally service their workstations.

- ♦ **Local application access:** Users need local access to their applications no matter where they connect to their network

You must manually create duplicate Application objects and create a site list in each copy of the Application object.

For more information on site lists, see “[Setting Up Site Lists](#)” in the *Novell ZENworks 7 Desktop Management Administration Guide*.

- ♦ **Server overload:** A server loaded with various application files can be over-worked to service all of its workstations

If you use only Desktop Management, you can configure Load Balancing (sharing the distribution workload between servers) in an Application object to address a server overload condition by having multiple servers being able to perform the same service. However, you would need to do a lot of manual work to use this feature.

- ♦ **Server redundancy:** If a server loaded with various application files goes down, its workstations cannot receive those applications

If you use only Desktop Management, you can configure Fault Tolerance (server redundancy) in an Application object to address the situation where a server goes down by having multiple backup servers listed in the Application object. However, you would need to do a lot of manual work to use this feature.

Server Management provides solutions to resolve these geographic and manual work issues. To see how, continue with “[Distributed Applications in Server Management](#)” on page 274.

Distributed Applications in Server Management

Server Management provides a Desktop Application Distribution that allows you to minimize your network traffic, local application access, server bandwidth, and redundancy issues with less effort on your part.

For example:

- ♦ **Network traffic:** Create a Desktop Application Distribution that contains your applications, then the Subscribers in each of your geographic areas create local copies of these applications. There, Novell Application Launcher can use these local applications to service the Subscriber server’s users and workstations.

- ♦ **Local application access:** After creating and sending a Desktop Application Distribution, link up site lists so that users who travel between geographic locations can have local access to their applications.

For information on how Server Management sets up site lists, see [Step 13 on page 295](#).

- ♦ **Server overload:** Through the Load Balancing feature, you can utilize multiple servers to service a large number of users or workstations via Novell Application Launcher. Simply use a

common working context for each of the servers receiving the Desktop Application Distribution. Then, you have multiple servers available for load balancing.

IMPORTANT: Load balancing is concerned with access to the source paths on Subscriber servers, not with access to the distributed Application objects.

- ♦ **Server redundancy:** Through the Fault Tolerance feature, you can have redundancy when servers go down by having other servers equally able to service your users and workstations via Novell Application Launcher. Simply use a common working context for each of the servers receiving the Desktop Application Distribution. Then, you have multiple servers available for fault tolerance.

IMPORTANT: Fault tolerance is concerned with access to the source paths on Subscriber servers, not with access to the distributed Application objects.

To do these things, you simply need to:

1. Create one Desktop Application Distribution for an application, or group of applications.
2. Send the Distribution to multiple servers.
Server Management automatically configures the application according to each server's environment.
3. Manually assign the necessary users or workstations to the groups that are associated with the new Application objects.
4. Click one button to link up the site lists.

Each server then has:

- ♦ Its own copy of an application's files on its file system
- ♦ Access to the Application object pointing to those files

The Application object is used to install the application on the workstations through Novell Application Launcher.

The Distribution process automatically does the multiple Application object creation, custom configuration, and file-copying work.

To further understand how Server Management can resolve these issues, continue with [Section 6.1.2, "Distributed Application Issues," on page 275](#).

6.1.2 Distributed Application Issues

When sending a Desktop Application Distribution, some content in an Application object is kept, some is not kept, and some is modified. The following sections explain this:

- ♦ ["Understanding Golden and Distributed Application Objects" on page 276](#)
- ♦ ["Maintaining a Golden Application's Attributes" on page 276](#)
- ♦ ["Maintaining Associations When Distributing Objects" on page 278](#)
- ♦ ["Maintaining Application File Rights" on page 279](#)
- ♦ ["Subscriber Working Context Conflicts" on page 280](#)
- ♦ ["Maintaining Source Paths" on page 281](#)

Understanding Golden and Distributed Application Objects

When you create a Desktop Application Distribution, you select an application object to be distributed. In Server Management, this is known as the “golden” Application object. All of the Application objects that are created by the Distribution are referred to as the “distributed” Application objects.

- ◆ [“Uniqueness of Golden Applications” on page 276](#)
- ◆ [“Synchronizing Golden and Distributed Applications” on page 276](#)

Uniqueness of Golden Applications

As an administrator, you should keep track of which objects are golden Application objects for the Distributions, because Application objects themselves do not have any visual designation in ConsoleOne® to identify them as such.

Because normal Desktop Management activity associated with Application objects can cause the object’s internal revision number to change, unnecessary deltas of a Distribution could be triggered and sent. For example, a Distribution rebuild could be triggered by a simple change in a User Group object that is associated with an application contained within the Distribution, which information is not even transferred to the distributed applications. Therefore, your golden Applications should not be used by users or workstations.

We recommend that you keep your golden Application objects in a unique Novell eDirectory™ context and associate users and workstations to only the distributed Application objects. For more information, see [Section 6.4, “Rebuilding Desktop Application Distributions,” on page 296](#).

Synchronizing Golden and Distributed Applications

When a Distribution is rebuilt and resent, all distributed Application objects are synchronized with the golden Application object. In other words, if you make important changes in a distributed Application object, but not the golden Application object, then you rebuild and send the Distribution again, you could lose your changes, because the Distribution only uses the content in the golden Application object to update the distributed objects. Therefore, the golden Application objects are the only objects that you should modify when you want to re-send the Distribution.

However, you can make changes to distributed Application objects that will not be overwritten, if those changes are in the attributes that are not normally overwritten by a re-sent Distribution. This is explained in the next section.

Continue with [“Maintaining a Golden Application’s Attributes” on page 276](#).

Maintaining a Golden Application’s Attributes

Server Management distributes most attributes that exist in a golden Application object, but not all of them. Therefore, various outcomes can occur for the attributes contained in distributed Application objects any time a Distribution is rebuilt.

The following sections provide information on when attributes are or are not distributed:

- ◆ [“Attributes Distributed” on page 277](#)
- ◆ [“Attributes Not Distributed” on page 277](#)
- ◆ [“Attributes Sent Only Once” on page 277](#)
- ◆ [“Attributes Modified” on page 277](#)

Attributes Distributed

If they can be modified in ConsoleOne, all attributes not listed in the following three sections are distributed as they exist in the golden Application object. These attributes are read from the golden Application object when building the Distribution and are sent every time Server Management creates or updates the distributed Application object.

All attributes contained in a golden Application object, not just the updated attributes, are updated in the distributed Application objects when a Distribution is rebuilt, sent, and extracted. This means that all distributed Application objects are kept in sync with their golden applications, except as noted in the next three sections.

Attributes Not Distributed

The attributes (listed by eDirectory attribute name in [Table 6-1](#)) are never read by the Distribution building process, and are not populated by Server Management in the distributed Application object:

Table 6-1 Location of Properties for Attributes Not Distributed

| Attribute Name | Location in the Application Object's Properties |
|----------------------|---|
| App:FS Rights Path | Common tab > File Rights subtab > Path column. |
| App:FS Rights Volume | Common tab > File Rights subtab > Volume column. |
| App:Printer Ports | Common tab > Drives/Ports subtab > Ports to be Captured list box. |

This list includes only those attributes that you can modify in ConsoleOne.

Attributes Sent Only Once

The attribute (listed by eDirectory attribute name in [Table 6-2](#)) is sent only once to provide an initial contact list:

Table 6-2 Location of Properties for Attributes Sent Only Once

| Attribute Name | Location in the Application Object's Properties |
|----------------|---|
| App:Contacts | Identification tab > Contacts subtab. |

This attribute is not updated by any subsequent Distribution updates. This prevents changes to this attribute in the distributed Application object from being overwritten by an original or updated contacts list in the golden Application object.

This attribute can be modified in ConsoleOne.

Attributes Modified

The attributes (listed by eDirectory attribute name in [Table 6-3](#)) are read from the golden Application object when the Distribution is built, but are modified to fit the Application object's new environment when the distributed Application object is created in the target server's working context:

Table 6-3 Location of Properties for Attributes Modified

| Attribute Name | Location in the Application Object's Properties |
|---------------------|--|
| ACL | NDS Rights tab > Trustees of This Object subtab. |
| App:Alt Back Link | Fault Tolerance tab > Remote Alternate App subtab. |
| App:Associations | Associations tab. |
| App:Back Link | Run Options tab > Application Dependencies subtab > Show Chain button. |
| App:Fault Tolerance | Fault Tolerance tab > Fault Tolerance subtab. |
| App:Load Balancing | Fault Tolerance tab > Load Balancing subtab. |
| App:Site List | Distribution tab > Link Up Site List button (which only displays if the Server Management snap-ins are installed in ConsoleOne). For how and why to use this button, see Step 13 on page 295 . |
| Application GUID | Distribution Options tab > Options subtab > GUID field. |
| creatorsName | Listed on the Other tab (you must click Show Read Only to view). |
| modifiersName | Other tab (you must click Show Read Only to view). |
| Object Class | Listed on the Other tab. |
| Revision | Listed on the Other tab (you must click Show Read Only to view). |
| Used By | Listed on the Other tab (you must click Show Read Only to view). |

This list includes only those attributes that you can modify in ConsoleOne, and they are only displayed in an Application object when needed to define the application.

Continue with [“Maintaining Associations When Distributing Objects” on page 278](#).

Maintaining Associations When Distributing Objects

When configuring a Desktop Application Distribution, you can specify to maintain associations. This means that you want attribute associations set in the golden Application object to be maintained in the distributed Application object that is created by the Distribution.

The Desktop Application Distribution requires some manual processes, such as adding the applicable users or workstations to the distributed Application object, which is empty of this information in Desktop Application Distribution object. This is because users and workstations can be different for each server receiving a distributed application.

If you select the Maintain Associations option, then attribute associations are handled in the following way:

- ◆ **Maintained:** User Group, Workstation Group, Organization, and Organizational Unit objects.

These are trusted groups and containers (within the source root container). They are maintained in the following manner:

- ◆ **Created new:** Group and container objects, if they do not exist.

You need to manually populate them with the users and workstations who need the distributed applications.

- ◆ **Not overwritten:** Group and container objects, if they already exist.

If group and container objects already exist and have been assigned to the distributed Application object, those settings are not overwritten, because they could already be populated with the users and workstations that need to use the distributed applications. If you want to add other users or workstations to existing groups or containers, you must add them manually.

- ◆ **Not created:** User and Workstation objects.

You can add the applicable users and workstations to the distributed Application objects after the Distribution has been extracted.

The Maintain Associations option is required when you distribute chained application information and folders. This is explained under [“Chained Applications in Distributions” on page 287](#).

Continue with [“Maintaining Application File Rights” on page 279](#).

Maintaining Application File Rights

File rights that you set in a golden Application object are not passed to the distributed Application objects, because file locations vary from server to server and cannot be anticipated.

The Desktop Application Distribution requires some manual processes, such as adding additional rights for file access. These processes are in addition to the minimums set by ZENworks when creating a distributed Application object. (The minimum rights might be enough for most applications.)

Review the following sections to understand how file rights are handled in Desktop Application Distributions:

- ◆ [“File Rights Are Not Distributed” on page 279](#)
- ◆ [“File Rights and Groups” on page 279](#)
- ◆ [“Chained Applications and File Rights” on page 280](#)
- ◆ [“Setting File Rights” on page 280](#)
- ◆ [“Setting Trustees and Shares Instead of File Rights” on page 280](#)

File Rights Are Not Distributed

File rights that are explicitly assigned in the Application object using the Rights to Files and Folders tab are not transferred, but are reset to the minimum necessary (Read and File Scan) for users to use the distributed applications. They are set when the distributed Application object is both created and then associated to a container or group.

File rights assigned in the Common > File Rights tab in the Application object are also not distributed.

You can later grant additional rights on these tabs in the distributed Application object and ZENworks does not remove or replace them.

File Rights and Groups

If a user or workstation is a member of a group that is distributed in the Desktop Application Distribution, then individual file rights for the user or workstation do not need to be set. The user or workstation obtains its rights to the application by virtue of its membership in the group.

Chained Applications and File Rights

If a chained application is used, all applications in the chain that require rights to a directory must be associated to a user or workstation group or a container in the golden Application's tree structure, because individual user or workstation objects' rights are not maintained in distributed Application objects.

Setting File Rights

To provide the Read and Write access rights to the files belonging to the chained application, in the Rights to Files and Folders tab in the User or Workstation Group object, assign the file rights.

Setting Trustees and Shares Instead of File Rights

Server Management does not set individual rights on files for NetWare-only trustees are set on the directories that contain the files, and rights are always Read and File Scan. Therefore, on NetWare servers you should grant users Read rights to the directory where the application's files are distributed. For example, if you have the files copied to the \apps directory, users would need Read rights to the \apps directory in order to use the application whose files were copied there.

Server Management also does not set file rights in Windows. Therefore, you should set up individual shares for users to have access to the application's distributed files.

Continue with [“Subscriber Working Context Conflicts” on page 280](#).

Subscriber Working Context Conflicts

Whether your Subscribers all use the same working context or a unique working context depends on your application distribution design needs. You might have all of the Subscribers who receive a Desktop Application Distribution use the same working context if you want load balancing or fault tolerance to be used. For more information, see [Section 6.1.1, “The Purpose of Desktop Application Distributions,” on page 273](#).

Where there are multiple Subscribers using the same working context, an eDirectory collision is possible. In other words, multiple Subscribers cannot extract their copies of the same Desktop Application Distribution at the same time to the same working context in the tree.

For example, if two Subscribers extract an application at the same moment and create an Application object in two different eDirectory replicas, this causes a problem in eDirectory synchronization. When eDirectory finds the two different objects, but with the same name and the same timestamp in the two different replicas, eDirectory resolves this by renaming one of the objects by appending a number to the collision object's name (for details, see TID 10062001 in the [Novell Support Knowledgebase \(http://support.novell.com/search/kb_index.jsp\)](http://support.novell.com/search/kb_index.jsp)).

If you use the same working context for a group of Subscribers, then you must make sure that each Subscriber's Extract schedule fires at a different time, allowing enough time between these schedules for extraction to be completed by a Subscriber before the next Subscriber begins extracting.

If you have each Subscriber use a unique working context, all Subscribers can then extract their copies of the same Desktop Application Distribution at the same time, and no eDirectory collisions occur.

If a Distribution is set to extract immediately, the same scenario can exist.

Continue with [“Maintaining Source Paths” on page 281](#).

Maintaining Source Paths

Many applications require supporting files, and the paths to those source files must be established in the Application objects in Desktop Management. This is known as the “source path.”

This section applies to Desktop Application Distributions containing Application objects that use source paths. For applications that require only an executable file (such as `notepad.exe`), source paths are not required in their Application objects.

Review the following sections to understand how source paths are used in Desktop Application Distributions:

- ◆ [“Source Path Usage in Server Management” on page 281](#)
- ◆ [“Purpose of the SOURCE_PATH Macro” on page 283](#)
- ◆ [“How the SOURCE_PATH Macro’s Values Are Interpreted” on page 284](#)
- ◆ [“MSI Applications and Source Paths” on page 284](#)
- ◆ [“Defining a Variable in Server Management” on page 284](#)
- ◆ [“Using the Source Path Option in the Distribution” on page 285](#)

Source Path Usage in Server Management

To show what happens with the attribute between the golden and distributed Application objects during the Desktop Application Distribution process, [Table 6-4](#) lists where you can find and configure source paths in ConsoleOne, their purposes, how these locations are populated, and their distribution status.

Table 6-4 *Source Path Usage Information*

| Source Path Name (Type) | Location in the Application Object | Purpose | How Populated | Distribution Information |
|-----------------------------------|--|---|---|---|
| SOURCE_PATH (macro ¹) | Distribution Options > Application Files > Source column | Provides path resolution from the SOURCE_PATH macro. | From the SOURCE_PATH macro. | This is distributed for each application file listed under the Name column that uses it. |
| SOURCE_PATH (macro) | Common > Macros > Name column | Defines a source path (in the Value column) to be used by the Application object. | From entries that you make on this page when the distributed Application object is created. | This is distributed with modifications to fit the Subscriber’s environment. If it is changed in the golden Application object, it is updated in the distributed Application object with the necessary modifications. This source path on the golden Application object should be kept stable, in order to avoid Novell Application Launcher distribution problems. |

| Source Path Name (Type) | Location in the Application Object | Purpose | How Populated | Distribution Information |
|---------------------------|------------------------------------|---|--|---|
| Package Source List (box) | Common > Sources | Provides a list of source paths for the Load Balancing and Fault Tolerance properties to use. | From the SOURCE_PATH macro, from each Subscriber using the same working context that receives and extracts the Distribution, and from any entries you make using the Add button. | <p>Only the SOURCE_PATH macro's entry is duplicated by Server Management using the long (DNS) version of the path.</p> <p>The listed source paths must be either valid UNC paths or variables that resolve to valid UNC paths.</p> <p>Each listed source path points to a complete set of the application's files that are located on the Distributor server's file system. (The Distributor cannot gather its Desktop Application Distribution's files from other servers.) These actual source files pointed to by the source paths are overwritten every time the Distribution is rebuilt and sent again.</p> <p>This field on the distributed Application object is cumulative, and is not overwritten when the Distribution is re-sent. Its entries come from selecting Load Balancing or Fault Tolerance for the Subscribers receiving the Distribution that use the same working context, or your use of the Add button on the distributed Application object. However, if you make a change to the SOURCE_PATH macro in the golden Application object, that source path is updated in this list box and is inserted first in the list. The previous source path is not replaced, but is left in the list. It is no longer valid, and you can delete it.</p> |

| Source Path Name (Type) | Location in the Application Object | Purpose | How Populated | Distribution Information |
|-------------------------|------------------------------------|--|--|---|
| Source List (box) | Fault Tolerance > Fault Tolerance | Provides a list of servers that can provide redundancy in case a server being used for Novell Application Launcher work goes down. All source paths listed must point to identical application file sets; otherwise, the distributed applications can fail to be created correctly. | From each Subscriber using the same working context that receives and extracts the Distribution. | This information is not distributed from the golden Application object. You must populate this field by sending the Distribution to multiple Subscribers that are using the same working context. |
| Source List (box) | Fault Tolerance > Load Balancing | Provides a list of servers that can provide load balancing among them for doing Novell Application Launcher work. All source paths listed must point to identical application file sets; otherwise, the distributed applications can fail to be created correctly. | From each Subscriber using the same working context that receives and extracts the Distribution. | This information is not distributed from the golden Application object. You must populate this field by sending the Distribution to multiple Subscribers that are using the same working context. |

¹ A “macro” in Desktop Management has the same functionality as a “variable” in Server Management.

Purpose of the SOURCE_PATH Macro

The SOURCE_PATH macro defines the source path in an Application object to where its application’s files reside on the Distributor server’s file system. This is where the Distributor accesses the application files for building the Desktop Application Distribution.

The SOURCE_PATH macro’s value is used to create the location on the Subscriber server’s file system where those application files are to be placed by the Subscriber when it creates the distributed Application object.

The information in the value of the SOURCE_PATH macro includes:

- ◆ Server identification (the Distributor server) in either the server name, IP address, or full DNS name
- ◆ Volume or drive on the Distributor server
- ◆ User-defined path information (if provided in the wizard)

- ◆ Application path information (selected in the wizard)

Some examples:

```
server1.novell.com\sys\apps\acrobat
```

```
server1.novell.com\n\apps\acrobat
```

This is resolved to a valid UNC path, such as:

```
\\server1\sys\apps\acrobat
```

If you include a macro (or variable) within the value of the SOURCE_PATH macro, Server Management does not resolve that embedded information. Server Management only resolves the SOURCE_PATH macro's value to a valid UNC path.

Mapped drive letters can also be used if a global policy variable is defined.

How the SOURCE_PATH Macro's Values Are Interpreted

Tiered Electronic Distribution searches variables to find a match with the golden Application object's source path. For example, if the source path in the golden Application object is `n:\apps\acrobat`, the following order is searched to find a match:

```
n:\
N:\
n:
N:
n
N
```

However, if the golden Application object's source path is `N:\apps\acrobat`, the following order is searched to find a match:

```
N:\
n:\
N:
n:
N
n
```

MSI Applications and Source Paths

The `.mst` files can be entered (on the MSI > Transforms tab) without specifying the file's full path, because Server Management uses the source path defined for the Application object to find these files when building the Desktop Application Distribution. However, this is only true when the `.mst` files are in the same directory as the MSI file.

Defining a Variable in Server Management

Historically, mapped drives have been embedded into an Application object as a means of launching that application from a mapped drive on the desktop. Server Management uses variables to distribute applications that use drive mappings.

Because volume names and mapped drives for the Distributor and all of the Subscribers receiving the Distribution can be different, variables allow you to identify these locations with a value that is interpreted by the Distributor and each Subscriber.

You can define variables globally and individually:

- ◆ **Globally using the Tiered Electronic Distribution policy**

Variables defined in the Tiered Electronic Distribution policy (Service Location Package) are available to all Subscriber objects associated with the policy, such as associating the policy package to the parent containers of the Subscriber objects. For the policy to be in effect for each Subscriber, make sure on the Variables property page that the Include Policy check box is selected.

The policy package must also be associated with the parent container of the Distributor object. The variable definition in the policy ensures that the Distributor knows where to gather the application files from.

If both the Distributor and Subscribers use the same variable value, then only one Tiered Electronic Distribution policy is needed, and you can associate its Service Location Package to the parent containers of both the Distributor and the Subscribers.

For example, the mapped drive source path for a golden Application object is `n:\applications\acrobat` and you want `n:` to represent the `sys:\public` directory on the Distributor server. To create the variable, in the Tiered Electronic Distribution policy select the Variables tab, then enter the `n:` for the variable and `sys:\public` for its value. Then, the Distributor can find the `\applications\acrobat` directory on its `sys:` volume when it needs to build the Distribution.

For more information on this policy, see [“Tiered Electronic Distribution” on page 209](#).

- ◆ **Individually in each target Subscriber’s properties**

You can define a variable for any Subscriber object. This definition overrides the same variable if it is defined in a Tiered Electronic Distribution policy that the Subscriber is associated with.

This is useful for when the Subscriber server’s volume name or mapped drive is different than the Distributor server’s (so they can’t use the same Tiered Electronic Distribution policy), or you have a variety in volume names or mapped drives among the Subscribers receiving the Distribution.

For information on how to define variables on Subscribers, see [Section 9.3, “Defining a Variable,” on page 346](#).

Using the Source Path Option in the Distribution

If a golden Application object uses mapped drives, enable the Keep the Same Source Paths for the Replicated Objects option when running the Desktop Application Distribution Wizard. Enabling this option causes the Distribution to retain golden application source paths for when a mapped drive designation must be used by the application that is distributed. The value of the mapped drive determines where the application files are copied.

If the golden application source paths are mapped drives, but you want the distributed applications to use a UNC path according to the extraction directory, then you do not need to select this option, but you must define the Tiered Electronic Distribution policy in the Service Location Package with variables defined for the mapped drives. This package must be associated to the Distributor with the variable defined in order for the Distribution build to work. It should also be associated with containers for the Subscriber objects, or any container above them, if they have the same drive mappings.

Key points about this option:

- ◆ If a golden Application object's Package Source List box contains a mapped drive, you must enable the Keep the Same Source Paths for the Replicated Objects option. The mapped drive letter is treated like a variable that needs to be resolved on both the Distributor and Subscriber to complete the valid UNC path.
- ◆ If a golden Application object's Package Source List box contains a drive mapping that is local to a server other than the Distributor server, no application files can be gathered or distributed, because all files to be included in the Distribution must be contained on the Distributor server's file system.
- ◆ Enabling this option affects all Application objects in the Distribution, including chained applications. Therefore, all mapped drive properties for each of the Application objects included in the Distribution are distributed to keep their golden application source paths, and each application and chained application must have mapped drives for the source path.
- ◆ If you select this option, only the Application object's Default Directory Path is used, because the Application Destination Directory Path field in the next wizard page is disabled. Therefore, you cannot change the path.
- ◆ When you select this option, or if you leave it unselected, that choice becomes the permanent use of this option for the Distribution. This is done to prevent problems that can occur from alternating between using and not using an Application object's mapped drives.
- ◆ For chained applications, source paths are treated the same for all chained applications as they are for the first application that the others are chained to.

6.1.3 Miscellaneous Issues

- ◆ [“Application Dependencies and Requirements” on page 286](#)
- ◆ [“Chained Applications in Distributions” on page 287](#)
- ◆ [“Extended Characters in Directory Paths” on page 287](#)
- ◆ [“Tree to Tree Distributions” on page 288](#)
- ◆ [“Site Distribution Objects” on page 288](#)

Application Dependencies and Requirements

Dependencies and requirements can be confusing with regard to the distribution of attributes:

- ◆ **Dependency:** An application dependency, such as a chained application, is updated in a distributed Application object when the Distribution is rebuilt, sent, and extracted.

To view application dependencies: in the properties of an Application object, click Run Options > Application Dependencies.

- ◆ **System requirement:** A system requirement, such as an operating system for the application to run on, is updated in a distributed Application object when the Distribution is rebuilt, sent, and extracted.

To view system requirements: in the properties of an Application object, click Availability > Distribution Rules.

One exception is that an application requirement is not updated in a distributed Application object when the Distribution is rebuilt, sent, and extracted. Instead, we recommend using application dependencies.

Chained Applications in Distributions

If multiple applications contain the same chained application, the application's files are only contained once in the Distribution. This reduces the Distribution's file size.

For example, if you distribute several icons (each its own Application object) that each require an office software suite, that suite software is only included once in the Distribution.

If your Desktop Application Distribution has chained applications, you must enable the Maintain Associations option when configuring the Distribution.

Chained applications in Distributions are only available in ZENworks for Desktops 4.x and later.

For more information on understanding and setting up chained applications, see “[Advanced Distribution: Configuring Application Dependencies and Chains](#)” in the *Novell ZENworks 7 Desktop Management Administration Guide*.

Extended Characters in Directory Paths

If extended characters (such as ê, ë, ì, or í) exist in the path to the `.fil` files for an AOT Application object, you must define a code page variable for both the Distributor and its Subscribers.

The code page variable is necessary for the Distributor so that it can gather the applications files from its file system when it builds the Distribution, and it is necessary for the Subscribers so that they can successfully copy the application files from the Distribution to their file systems. In other words, the code pages used by the Distributor and Subscribers must contain the extended characters used in the paths contained in the Distribution.

To create the code page variable:

- 1 Determine the code page used by ConsoleOne for the international characters to be used in the Distribution.

The code page must come from the workstation used to create the golden Application objects that have extended characters in the paths to their AOT files.

You can use the `encoding.cmd` utility included in the `\codepageutility` directory on the *ZENworks 7 with Support Pack 1 Companion 2* CD to determine the necessary code page. Instructions on how to use this utility are contained in the `readme.doc` file included in the `\codepageutility` directory.

- 2 In ConsoleOne, create a Service Location Package and access its properties.
- 3 Click the *Tiered Electronic Distribution* policy to enable it, then click *Properties*.
- 4 Add the following variable in the policy: `CODE_PAGE`.
- 5 Enter the desired code page as the variable's value.

For example, Cp1252.

Code page values are case sensitive.

- 6 Click *OK* to save the policy.
- 7 Associate the policy package to the container of the Distributor object.

- 8 To provide Subscribers access to the code page, do one of the following:
 - ♦ Associate the policy package to the container of the Subscriber objects receiving the Distribution.

This causes the variable to be available for use by the Subscribers, providing access to the code page.
 - ♦ Define the same code page variable (as in [Step 4](#) and [Step 5](#)) in each Subscriber object's properties.

This provides Subscribers access to the code page.
- 9 Rebuild the Distribution, if it has already been created.

Tree to Tree Distributions

You can send Desktop Application Distributions to other trees. However, ZENworks 7 Distributions cannot be sent to ZENworks for Servers 3.0.2 Subscribers because of new schema extensions for ZENworks 7 and later.

Site Distribution Objects

ZENworks 7 Server Management does not use a Site Distribution object. Previous versions of ZENworks might have used Site Distribution objects with this Distribution type.

6.2 Requirements

The following requirements must be met before creating and sending Desktop Application Distributions using Tiered Electronic Distribution:

- In order to use Novell Application Management with ZENworks 7 Server Management, you must have ZENworks for Desktops 4.01 or ZENworks 7 Desktop Management or later installed. Chained applications in Desktop Application Distributions are only supported in ZENworks for Desktops 4.01 and later.
- Desktop Management and Server Management must both be installed to the same tree, including their respective schema extensions.
- For golden Application objects to be functional in a Desktop Application Distribution, the snap-ins for both Server Management and Desktop Management must be installed in ConsoleOne.
- Make sure all of the associations in the golden Application object are in the source root context or below.

IMPORTANT: If even one of your associations is outside the source root context, the Distributor fails to build the Distribution.

- For Windows and Linux servers, you must have a shared location established for extracting the Distribution's files, where all users can have access to those files.
- The source path must point to application files that are located on the Distributor server's file system, because the Distributor cannot gather files from other servers' file systems.

If the Package Source List box contains a local drive mapping, no application files are gathered or distributed.

If the Package Source List box contains a mapped drive, Keep the Same Source Paths for Replicated Objects must be selected. The drive letter is treated like a variable that needs to be resolved on both the Distributor and Subscriber to complete a valid UNC path.

Use a policy to define the variables on a Distributor. On the Subscriber you can use either the variable list in the Subscriber object, or a policy that is associated to the container where the Subscriber object resides.

- ❑ The Subscriber object must have the Working Context attribute defined. This is the eDirectory context where the Subscriber creates the objects related to the Desktop Application Distributions that it receives.

Multiple Subscribers can use the same working context if you intend to use them for load balancing or fault tolerance.

- ❑ If extended characters (such as ê, ë, ì, or í) exist in the path to the `.fil` files for an AOT Application object, you must define a code page variable for the Distributor and Subscribers. For more information, see [“Extended Characters in Directory Paths” on page 287](#).
- ❑ Determine whether you are using Samba or NCP shares for client access to files on Linux or OES Linux servers, then set up that access. For more information, see [Appendix G, “Client Access in Linux,” on page 425](#).

6.3 Creating a Desktop Application Distribution

- ♦ [Section 6.3.1, “Understanding the Desktop Application Distribution Wizard,” on page 289](#)
- ♦ [Section 6.3.2, “Creating the Distribution,” on page 290](#)

6.3.1 Understanding the Desktop Application Distribution Wizard

ZENworks uses Tiered Electronic Distribution to distribute Application objects to other locations in the same tree or other trees. Using a Desktop Application Distribution, the original files associated with the applications are copied to the appropriate server locations where they can be used to locally service user groups and workstation groups associated with the distributed Application objects.

To distribute applications, you use the Server Management Desktop Application Distribution Wizard to configure the Distribution. This includes:

- ♦ Determining the destination’s tree context
- ♦ Determining whether to maintain the associations between user/workstation groups or containers and the applications
- ♦ Determining whether to have automated load balancing or fault tolerance
- ♦ Determining how to trigger rebuilds of the Distribution
- ♦ Selecting the applications
- ♦ Determining the file copying paths

To create a Desktop Application Distribution using the wizard, continue with [Section 6.3.2, “Creating the Distribution,” on page 290](#).

6.3.2 Creating the Distribution

IMPORTANT: You can also perform these step in iManager instead of running this wizard in ConsoleOne. Instructions are contained in the context-sensitive help in iManager.

- 1 Fulfill all of the requirements listed under [Section 6.2, “Requirements,” on page 288](#), including creation of a code page variable if necessary.
- 2 In ConsoleOne, right-click the container where you want the Distribution object located, click *New*, click *Object*, select *TED Distribution*, then click *OK*.
- 3 Provide a name for the Distribution.

IMPORTANT: Periods (.) are not allowed in Distribution names. Instead, use dashes (-) or underscores (_) as word separators. If you use a period in the Distribution name, the Distribution is not sent, and the Distributor is not reloaded after it has been exited.

- 4 To give a Distributor ownership of the Distribution, browse and select the Distributor object, click *Define Additional Properties*, then click *OK*.

The Distribution object’s properties are displayed.

- 5 Select the *General* tab and fill in the *Settings* fields:

Active: Required. In order to make a Distribution available to Subscribers, it needs to be active.

Use digests: Digests are used by Distributors and Subscribers to verify that Distributions have not been tampered with while in transit. The digest provides an MD5 checksum for the Subscriber to compare.

Digests also detect corruption in a Distribution’s package. If corruption is present, the Subscriber renames the `distfile.ted` Distribution file to `distfile.corrupt` and the Distribution is rebuilt and sent the next time the Channel’s schedule fires.

Encrypt: You can have the Distribution encrypted if you are sending it across non-secured connections. Encryption provides security for the Distribution during transit between the Distributor and Subscriber when they are not within the same firewall. Select either Strong or Weak encryption.

You must also have NCI 2.6.4 or later installed to each of these servers for encryption to work (see [“Installing NCI 2.6.4” on page 54](#)). Older versions of NCI are not compatible with version 2.6.4 or later.

Maximum revisions: This number helps you to control disk space usage by determining how many versions of a particular Distribution are kept in the Distributors’ and Subscribers’ working directories. The default is 10. Select *Limited* and enter a number.

Increase the number if data is changing often and the changes are minimal (smaller delta files). Decrease the number if data is not changing very often, or if a significant amount of data is changing (larger delta files).

The following e-mail options are available if you set a maximum number. If you select *Unlimited*, these options are dimmed.

- ♦ **Approaching maximum revision email notification list:** Contains the e-mail addresses of anyone who is to be notified when a Distribution is approaching the maximum revisions set in the *Maximum revisions* field. Here, you can either remove a single or all displayed addresses.

- ♦ **Email address (maximum revision notification):** You can add e-mail addresses to the list in *Approaching maximum revision email notification list*. Just enter an e-mail address and click *Add* and it is displayed in the listing.
- ♦ **Send notifications when Distribution revision is ___ or less of reaching maximum revisions:** Enter a number to indicate how close “approaching” is. When the current revision number of Distribution plus this number equal the maximum revisions, an SMTP notification is sent to the listed addressees.
SMTP must be configured and its e-mail server address listed in the next field.
- ♦ **Email server address:** The SMTP server used to send the e-mail notifications. For example, mail.novell.com.
For information on configuring SMTP e-mail notifications, see “SMTP Host” on [page 208](#).

Priority: You can give the Distribution a priority that determines how it is sent in relation to other Distributions. A High priority means it is sent before Medium or Low priority Distributions.

Distributor: Displays the DN of the Distributor object that builds and sends this Distribution. You selected the Distributor when you created the Distribution object.

Description: Provide useful details about the Distribution, such as the name of the desktop application, the files and directories it contains, intended user groups, and so on.

6 Click *General > Restrictions*.

You can select whether to have platform restrictions for the Distribution itself. This is not a restriction for the distributed Application object.

No restrictions: This option is selected by default. To determine platform restrictions, select this option to disable it, then select the check boxes corresponding to the platforms you want to receive this Distribution.

Platforms with check boxes not selected cannot receive the Distribution. In other words, you restrict sending to a platform by deselecting the No Restrictions option and not selecting the platform.

The available options are:

- No Restrictions
- NetWare All
- NetWare 4.x (earlier versions of ZfS supported these platforms)
- NetWare 5.0 (earlier versions of ZfS supported this platform)
- NetWare 5.1
- NetWare 5.x (earlier versions of ZfS supported these platforms)
- NetWare 6.x
- Windows Server

No Restrictions means that the Distribution can be sent to any of these platforms.

If you select NetWare All, you do not need to select any of the individual NetWare platforms.

7 Select the *Type* tab, select Desktop Application in the Select Type drop-down box, then click *Setup*.

The Desktop Application Distribution Wizard is started. iManager provides its own interface in place of this wizard.

You can also start this wizard from the Desktop Application Agent properties page by clicking *Modify*.

7a Click *Next* after reading the introductory information.

7b Fill in the fields, then click *Next*.

Maintain source tree structure: Duplicates the source tree's structure at the destination's location (the target Subscriber's working context) for placing the distributed Application objects. If you are selecting chained applications, you must select this option.

For more information, see [“Maintaining Source Paths” on page 281](#).

Source root context: Select a container to be used as the root container for the golden Application objects to be distributed. You should select golden Application objects only from this root container and its subordinate containers.

Maintain associations: Distributes the associated groups or containers at the target location if they do not exist. However, users or workstations contained in the groups or containers in the source location are not distributed.

For more information, see [“Maintaining Associations When Distributing Objects” on page 278](#).

You must enable this option if you have chained applications in the Distribution. For more information, see [“Chained Applications in Distributions” on page 287](#).

IMPORTANT: Rights previously set in the associated user/workstation groups or containers that are maintained are not distributed, but set to the minimum necessary in the distributed groups or containers so that users can use the applications.

Overwrite Existing Target Folder Object Attributes: When selected, this check box causes existing target folder objects to be overwritten with the relevant content of the source folder objects, meaning all previous folder associations for the application are replaced by the new folder associations.

To retain previous folder associations while adding new folder associations, deselect this option. (This is the default function previous to ZENworks 7 for how folder objects are handled. The option to overwrite is added in ZENworks 7.)

Always replicate association flags: Causes the launch configuration flags for each group or container associated with the golden Application object to be replicated with each distributed application.

Load balance and fault tolerance support: Choose whether to use automated load balancing, fault tolerance, or neither:

- ◆ **Load balance:** Automates spreading server workloads over the servers being used for the Desktop Application Distributions. The functionality of fault tolerance (redundancy) is automatically accomplished through load balancing.
- ◆ **Fault tolerance:** Allows a server being used for Desktop Application Distributions to assume the distribution duties of another server that goes down. Fault Tolerance does not provide load balancing.
- ◆ **None:** Neither option is applied. You must individually configure each distributed Application object for load balancing or fault tolerance, if you want that support on an individual basis.

For these two features to work:

- ◆ Multiple Subscribers receiving the Distribution must be using the same working context

- ◆ The User Source List button must be selected on the Fault Tolerance > Fault Tolerance or the Fault Tolerance > Load Balancing properties pages of the Application object

Depending on the selected options, the Load Balancing or Fault Tolerance pages are populated with the file locations on all servers that share this working context.

For more information, see [“Distributed Applications in Server Management” on page 274](#).

Rebuild only if any application version number changes: Allows you to control Distribution rebuilding based on the Build schedule. Select this check box to withhold modifications to a golden Application object until you are ready to release them.

Regardless of the status of this check box, if applications are added to or removed from the Distribution, it is rebuilt according to its Build schedule.

For more information, see [Section 6.4.2, “Triggering a Rebuild,” on page 296](#).

7c Click *Add* to browse for and select golden Application objects.

Do not browse above the root directory that you established in the previous wizard page, especially if you have selected the Maintain Source Tree Structure option.

IMPORTANT: The Desktop Application source files must reside in the Distributor server’s file system. The Distribution cannot be gathered from another server’s file system.

7d Select the *Keep the same source paths for the replicated objects* option if you want to retain the golden Application object’s source path when the mapped drive feature is used in the distributed application.

For more information, see [“Maintaining Source Paths” on page 281](#).

7e Provide the destination volume or shared folder.

The application files distributed are those that are associated with the golden Application objects you selected in the previous wizard page.

You can provide a variable instead. If you use a variable, it must be defined in the destination Subscriber server’s properties to point to the target server’s volume or shared folder.

This volume (NetWare) or shared folder (Windows and Linux) becomes the root location for placing subordinate directories where the application files are copied.

7f To use only an application’s default path, click *Application’s default directory path*, which is placed beginning with the root location you specified in [Step 7e](#).

or

To provide a user-defined directory path to the application’s files, click User-Defined Directory Path, then provide your path information.

The path you specify is used in the following manner:

- ◆ The volume or shared folder name remains unchanged (as specified in [Step 7e](#)).
- ◆ Your path information is inserted after the volume or shared folder name.
- ◆ Part of the application’s default path is appended to your path information, beginning with the default path’s immediate parent directory to the application’s files. Any default path information that was above the immediate parent directory is replaced by your path entry.

The result is a customized directory path that begins with the volume or shared folder, has your user-defined path information next, and ends with the application's immediate directory. For example, suppose the default path to the application's executable file (`application.exe`) might be:

```
\application_root_directory\application_subdirectory
```

and you enter `\mypath` for your user-defined path; the new full path to the executable is now:

```
c:\mypath\application_subdirectory\application.exe
```

If you entered `C:` as the shared folder and `\mypath` as your user-defined path, `\application_root_directory` is replaced by `\mypath`, and `\application_subdirectory` is the immediate parent directory to `application.exe`.

7g Click *Next* to continue.

The Summary page is displayed.

7h To make changes, click *Back*.

7i When you have finished configuring the Distribution object, click *Finish* to exit the wizard.

You can edit the Distribution at any time on the *Type* tab of the Distribution object by clicking *Modify*.

8 Select the *Channels* tab, click *Add*, then browse for and select the Channel for this Distribution. Each Distribution must be associated with at least one Channel if it is going to be used to push data to a Subscriber. A Distribution is sent to all Subscribers that are subscribed to the selected Channel.

9 Select the *Schedule* tab, then select a Build schedule:

[Section B.1, "Daily," on page 400](#)

[Section B.3, "Interval," on page 400](#)

[Section B.4, "Monthly," on page 401](#)

[Section B.5, "Never," on page 401](#)

[Section B.8, "Run Immediately," on page 402](#)

[Section B.9, "Time," on page 402](#)

[Section B.11, "Yearly," on page 403](#)

10 Click *Apply* to create the Distribution.

You are prompted to copy additional security certificates.

11 Select *Yes* to resolve the certificates.

This copies the security certificates from the Distributor to Subscriber subscribed to the Channel.

For information, see [Section 7.1.6, "Resolving Certificates," on page 303](#).

12 Click *OK* to close the Distribution object.

The next time the Distributor reads eDirectory (this schedule is set in the Distributor object's properties), it retrieves all of the information about the new Desktop Application Distribution, such as Distribution details, the Build schedule, and so on.

The Distribution is built according to the Build schedule, sent according to the schedule set in the Channel object, and extracted according to schedule set in the Subscriber object.

If the Distributor throws an exception during the file gathering process, the Distribution is not built. The Distributor logs the failure in the reporting database.

If the Subscriber throws an exception during extraction, the process is not completed. The Distributor receives this information from the Subscriber and logs the failure in the reporting database.

After extraction, Desktop Management users whose objects are located in the associated containers, or are members of a distributed group, will have access to the desktop applications that were distributed.

IMPORTANT: For Desktop Application Distributions, a built-in delay exists to accommodate directory synchronization when you have multiple applications being distributed at the same time (whether by one or multiple Distributions).

Subscribers can receive Desktop Application Distributions all at the same time, but extract them one at a time. And, when there are multiple applications contained in one Distribution, the Subscriber creates the distributed Application object and copies the files one application at a time. The built-in delay helps directory synchronization for the newly-created Application objects to occur smoothly.

To determine how much additional time this built-in delay might add to the distribution process, multiply each application contained in a Desktop Application Distribution by 30 seconds.

As a rule of thumb, if an application being distributed includes multiple versions, such as one baseline and two deltas, each of these three versions receives the same 30-second delay. For example, if you are sending 10 desktop applications, and each has three versions, the completion of the Distribution extractions could take at least 15 minutes.

13 For Desktop Management users and workstations to have automatic access to their applications from any geographic location, you must link up the site lists:

13a Wait for the Desktop Application Distribution to be distributed and extracted by each Subscriber server that received it, because the distributed Application objects must be created and the application's files installed before you can link up the site lists.

13b In ConsoleOne, right-click the golden Application object that was used to build the Distribution, then click *Properties*.

13c Click the *Distributions* tab, then click the *Link Up Site Lists* button.

All distributed Application objects that were created from the golden Application object are displayed in the *Replicated Applications* list box, and all Distributions containing the distributed Application objects are listed in the *Distributions Currently In* list box.

The *Link Up Site List* button does the following:

13c1 For the golden Application object, it searches for each distributed Application object that was created from the Distribution and lists the full DN of those objects in the golden Application object's properties (in the *Replicated Applications* list box).

13c2 For each distributed Application object, it searches for the other distributed Application objects that were created from the Distribution; for the golden Application object, it lists the full DN of all of these objects in the distributed Application object's properties (in the *Replicated Applications* list box).

13c3 **Step 13c1** and **Step 13c2** are repeated for each distributed Application object.

13c4 For each Application object listed in the Replicated Applications list box, any Distributions associated with those objects are listed in the Distributions Currently In list box in each of these Application objects.

Thus, the golden Application object and all distributed Application objects have each other listed in their Replicated Applications list box, which allows users to have local access to the same application no matter where they connect to their network.

13d Click *OK* to close the golden Application object's properties.

13e Repeat [Step 13b](#) through [Step 13d](#) for each golden Application object that was used to build the Desktop Application Distribution.

You need to perform the site list link-up only on the golden Application objects.

6.4 Rebuilding Desktop Application Distributions

The following sections explain the different issues with rebuilding Desktop Application Distributions, including how to trigger a rebuild:

- ◆ [Section 6.4.1, "All Attributes Are Updated," on page 296](#)
- ◆ [Section 6.4.2, "Triggering a Rebuild," on page 296](#)

6.4.1 All Attributes Are Updated

All attributes contained in a golden Application object, not just the modified attributes, are updated in the distributed Application objects when a Distribution is rebuilt, sent, and extracted. This means that if you make a change to an attribute in a distributed Application object, such as a source path, that source path is overwritten by the source path data in the golden Application object. In other words, all distributed Application objects are kept in sync with their golden Application object. Exceptions to this are described in ["Maintaining a Golden Application's Attributes" on page 276](#).

A rebuilt Desktop Application Distribution includes all file changes made after the last time the Distribution was built.

6.4.2 Triggering a Rebuild

You can control when a Distribution is rebuilt by whether you select the Rebuild Only If Any Application Number Changes check box in the Desktop Application Distribution Wizard:

- ◆ ["Selecting the Check Box" on page 296](#)
- ◆ ["Leaving the Check Box Disabled" on page 297](#)

Selecting the Check Box

This feature is useful for withholding modifications to a golden Application object until you are ready to release them.

The Distribution is rebuilt according to its established Build schedule, but only after you have manually incremented the Version Number field in the golden Application object, or its dependent application, and the Distributor has read eDirectory to discover the Version Number field change.

If there are multiple applications in a Distribution, a version number change in only one of them triggers a rebuild of the Distribution for all of them.

The Version Number field is on the Distribution Options > Options tab of the Application object's properties.

Regardless of the status of this check box, it is rebuilt according to its Build schedule if applications are added to or removed from the Distribution.

Leaving the Check Box Disabled

If you do not select this option (it is unchecked by default), the Distribution is rebuilt according to its established Build schedule. In this case, there can be two scenarios:

- ♦ [“Modifying an Object” on page 297](#)
- ♦ [“Removing a Distributed Application Object” on page 297](#)

Modifying an Object

When you modify a Distribution object or one of its golden Application objects, its internal revision number is automatically changed, which triggers a rebuild of the Distribution according to its established Build schedule.

Modifications include adding or removing applications from the Distribution. However, if you simply update, add, or remove application files in the Distributor server's file system, this does not alter the internal revision number of the Desktop Application Distribution object. The ZENworks file synchronization feature does not apply to the files in Application objects. Therefore, no rebuild is triggered.

If you add, remove, or update any files belonging to a golden Application object, those changes are included when the next rebuild is triggered.

Removing a Distributed Application Object

Removing a distributed Application object causes a backlink to the golden Application object to change without any other changes being made to the object. This causes the internal revision number to change on the golden Application object, which triggers a rebuild of its Distribution according to the established Build schedule.

6.5 Cleaning Up Desktop Application Distribution Files

Tiered Electronic Distribution is not designed to clean up Desktop Management files, for example, when you might delete a golden Application object in Desktop Management. The distributed Application object created by the Desktop Application Distribution is not automatically deleted. You must manually remove this eDirectory object and the related application files on the server's file system.

To do this, search in ConsoleOne for the Application object and note its filename before deleting the golden Application object. This can make manually cleaning up easier after deleting a golden application.

You can verify that the distributed application exists after you've deleted the golden Application object and its files, and then remove the distributed version:

- 1 In ConsoleOne, delete the golden Application object.
- 2 In a file browser, delete the files related to the application.
- 3 In ConsoleOne, right-click the Distributor object associated with the golden Application object's Distribution.
- 4 Click *Refresh Distributor* from the drop-down menu, then click *Yes > OK*.
- 5 Check the destination location for the distributed Application object.
It should still be present, even though it no longer exists in the Distribution.
- 6 In ConsoleOne, delete the distributed Application object.
- 7 In a file browser, delete the files related to the distributed application.

6.6 Sending Desktop Application Distributions Tree-To-Tree

Desktop Application Distributions can be sent between trees. However, you must do the following for this to work:

- 1 Create an External Subscriber object in the Distributor's tree that points to the target server in the other tree where you want to send the Desktop Application Distribution.

This enables the Distributor server to send the Distribution directly to the target server using the IP address listed in the External Subscriber object.
- 2 Make sure the target server that is to receive the Desktop Application Distribution has a Subscriber object in its own tree, so that it has the rights to eDirectory for creating the distributed Application object in that tree.

This Subscriber must be within a working Tiered Electronic Distribution system.
- 3 Set the working context in the Subscriber object for the target server, if this was not done during installation.

If the working context is not set for the target server, authentication fails during the extraction process.
- 4 Create the Desktop Application Distribution (see [Section 6.3, "Creating a Desktop Application Distribution,"](#) on page 289).

Defining the Desktop Application Distribution is the same process, whether it is being sent within a tree or across trees.
- 5 Add the External Subscriber object to the Channel where the Desktop Application Distribution is listed.
- 6 Manually copy the certificates to the Subscriber in the target tree.
or
If you have a mapped drive, browse to the correct path when prompted.
- 7 Send the Distribution.

Security in Policy and Distribution Services

7

Novell® ZENworks® Server Management provides the following types of security for Policy and Distribution Services:

- ◆ [Section 7.1, “Distribution Security Using Signed Certificates and Digests,” on page 299](#)
- ◆ [Section 7.2, “Distribution Security Using Encryption,” on page 309](#)
- ◆ [Section 7.3, “Security for Inter-Server Communication Across Non-Secured Connections,” on page 313](#)

7.1 Distribution Security Using Signed Certificates and Digests

There are two features of Tiered Electronic Distribution that deal with security:

- ◆ **Certificates (required):** Security certificates are issued by each Distributor to all Subscribers receiving its Distributions to validate whether Distributions are from a trusted source, or have been tampered with. This security is automatically used by Policy and Distribution Services for all Distributions. However, there are actions you might need to take to get Policy and Distribution Services to create and process the certificates.

In order for a Subscriber to accept its first Distribution from a Distributor, it must have a certificate from that Distributor in its `\security` directory. After receiving its first Distribution from the Distributor, the certificate is first stored in the `.keystore` file, then the certificate is deleted from the `\security` directory.

The `.keystore` file is a repository of signed certificates from the various Distributors who send Distributions to the Subscriber. In other words, it provides the Subscriber with an accumulation of trusted sources for its Distributions.

You can view the content of the `.keystore` file in Novell iManager.

For information on security certificates for encrypted Distributions, see [Section 7.2, “Distribution Security Using Encryption,” on page 309](#).

- ◆ **Digests (optional):** You can have digests created for each Distribution at the time it is built. The digest provides an MD5 checksum for the Subscriber to compare against to determine whether a Distribution has been tampered with after it left the Distributor.

Digests also detect corruption in a Distribution’s package. In the case of corruption, the Subscriber renames the `distfile.ted` Distribution file to `distfile.corrupt` and the Distribution is rebuilt and sent the next time the Channel’s schedule fires.

The following sections provide more information on understanding, creating, and using certificates and digests:

- ◆ [Section 7.1.1, “Understanding Digests,” on page 300](#)
- ◆ [Section 7.1.2, “Understanding Certificate Usage in Policy and Distribution Services,” on page 300](#)

- ◆ Section 7.1.3, “Important Points about Certificates,” on page 301
- ◆ Section 7.1.4, “ConsoleOne User Rights and Certificate Copying,” on page 302
- ◆ Section 7.1.5, “Certificate File Locations,” on page 303
- ◆ Section 7.1.6, “Resolving Certificates,” on page 303
- ◆ Section 7.1.7, “Handling Invalid Certificates,” on page 304
- ◆ Section 7.1.8, “Certificate and Private Key Directories,” on page 307
- ◆ Section 7.1.9, “Creating Security Certificates for Non-Encrypted Distributions,” on page 307
- ◆ Section 7.1.10, “Manually Copying Certificates for Non-Encrypted Distributions,” on page 308

7.1.1 Understanding Digests

Important points about digests:

- ◆ Digests can be created for each Distribution at the time it is built. The digest is used by the Subscriber to determine whether a Distribution has been tampered with after it left the Distributor.
- ◆ Digests detects corruption in a Distribution’s package. In the case of corruption, the Subscriber renames the `distfile.ted` Distribution file to `distfile.corrupt` and the Distribution is rebuilt and sent the next time the Channel’s schedule fires.
- ◆ The Digest option is available for all Distribution types. The Digest check box is displayed on the General tab of the Distribution object’s properties.
- ◆ A digest adds to the build time. Factors that can affect build time using digests are CPU and hard drive speeds, amount of RAM, server workload, and so on.

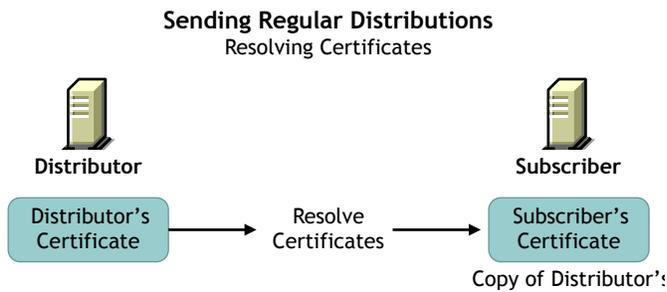
7.1.2 Understanding Certificate Usage in Policy and Distribution Services

A certificate is a security mechanism used by Policy and Distribution Services to ensure that the Distribution received by a Subscriber was actually sent by the Distributor owning that Distribution. Because configuration information can also be sent to the Subscriber, it ensures that the configuration information has been sent from a known Distributor and that the data has not changed.

All Subscribers must receive a valid security certificate from each Distributor that sends Distributions to them. Without a matching certificate, a Subscriber cannot receive Distributions from the Distributor.

Figure 7-1 illustrates the process of using certificates with Distributions:

Figure 7-1 Resolving Certificates

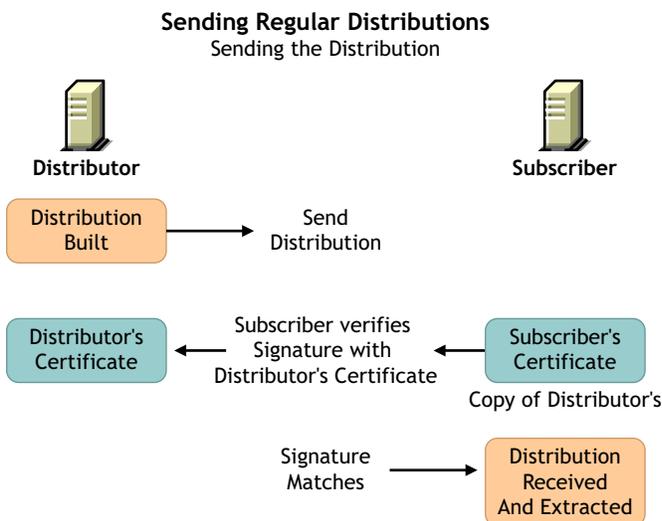


Before a Distribution is sent, certificates must be resolved. This ensures that the Distribution received by a Subscriber was actually sent by the Distributor owning that Distribution.

For information on resolving certificates, see [Section 7.1.6, "Resolving Certificates,"](#) on page 303.

After certificates have been resolved, the following illustrates how the Subscriber uses the certificate to ensure it is receiving a valid Distribution, as illustrated in [Figure 7-2:](#)

Figure 7-2 Sending the Distribution



7.1.3 Important Points about Certificates

- ◆ Certificates are issued by each Distributor to all Subscribers receiving Distributions from that Distributor. In order for a Subscriber to accept Distributions from a Distributor, it must have received a certificate from that Distributor.
- ◆ For security, certificate key pairs are created by the Distributor.
- ◆ The public key is written to the Distributor server's file system, which self-signs a certificate and stores it in Novell eDirectory™.
- ◆ The private key is stored in the Distributor object's properties and is used for encryption.
- ◆ The Subscriber software does not need to be running on the Subscriber server to have certificates copied to the server.

- ◆ The association of Distributions (owned by a Distributor) and Subscribers to a Channel determines which Subscribers should receive certificates from which Distributors.
 - ◆ A Distributor sends certificates to all Subscribers that subscribe to Channels where the Distributor has Distributions.
 - ◆ A Subscriber requests certificates from all Distributors that have Distributions in Channels to which it subscribes.
- ◆ A certificate can be passed from a Distributor to a Subscriber under the following circumstances:
 - ◆ When a Subscriber is initially subscribed to a Channel and you click OK to apply the changes.
 - ◆ When you right-click a Subscriber Object and select Resolve Certificates. The Subscriber then requests certificates from all Distributors that it receives Distributions from.
 - ◆ When a Distribution is listed in a Channel and you click OK to apply the changes.
 - ◆ When you right-click a Distributor Object and select Resolve Certificates. The Distributor sends certificates to all Subscribers that it sends Distributions to.

For information on resolving certificates, see [Section 7.1.6, “Resolving Certificates,” on page 303](#).

- ◆ When you add a Distribution or a Subscriber to a Channel. When you click OK, the Resolve Certificates? dialog box is displayed. If you answer Yes, certificates are sent by all Distributors who have Distributions associated with that Channel to all Subscribers subscribed to that channel.
- ◆ Manually copying a certificate file to a transfer medium (such as a diskette or local drive), then to the `\zenworks\pds\td\security` directory on a server.

Basically, any time the relationship changes between the Subscribers, Channels, or Distributions, a certificate can be passed.

- ◆ If a Distributor object is deleted and re-created to point to the same server, all certificates on the subordinate Subscribers become invalid. Certificates must be deleted from the Subscriber’s `\security` directory, then the Distributor must send the new certificates to those Subscribers.
- ◆ ConsoleOne copies the certificate files to Subscriber servers. Therefore, the client software on the workstation running ConsoleOne must have access to the Subscriber servers’ file systems. For Windows Subscriber servers, the Domain and Workgroup rights on the workstation must be set up to facilitate automatic certificate copying. Otherwise, a 1204a error is given.

7.1.4 ConsoleOne User Rights and Certificate Copying

The administrator using ConsoleOne® must have sufficient rights to the Subscriber server in order for a certificate to be copied to that server when the administrator resolves certificates in ConsoleOne. This is because when you use ConsoleOne to configure a Subscriber object to receive the Distributions from a particular Channel, the Distributors owning the Distributions in that Channel must send certificates to the Subscriber’s server.

For NetWare® Subscribers, the ConsoleOne user automatically has sufficient rights by virtue of being able to configure the Subscriber object.

For Windows Subscribers, administrator rights for the ConsoleOne user must be set up in Windows by selecting Active Directory Users and Computers, or selecting Local Users and Groups.

7.1.5 Certificate File Locations

Certificates are stored in the `\zenworks\pds\ted\security` directory on NetWare and Windows Subscriber servers, or in the `/var/opt/novell/zenworks/zfs/pds/ted/security` directory on Linux and Solaris servers.

WARNING: Make sure the `\security` directory is a non-public directory. This directory should not be read by anyone other than an administrator. The `.keystore` file is in the `\security\private` directory and is by default hidden from non-administrative users.

Certificates are usually named after the fully qualified DNS name of the Distributor server, such as `Distributor_Server001.Distributions.ZENworks.Novell.com.cer` or `Distributor_Server001.Distributions.ZENworks.Novell.com.csr`. The TCP/IP address of the server would be used for `.csr` files if a DNS name could not be resolved. The certificate would then be named using its IP address, such as `155.55.155.55.csr`.

7.1.6 Resolving Certificates

IMPORTANT: ConsoleOne copies the certificate files to Subscriber servers. Therefore, the client software on the workstation running ConsoleOne must have access to the Subscriber servers' file systems. For Windows Subscriber servers, the Domain and Workgroup rights on the workstation must be set up to facilitate automatic certificate copying. Otherwise, a 1204a error is given.

When you are automatically presented with the option in ConsoleOne to resolve certificates, determine the following to know whether to select Yes or No:

- ◆ If the Distributor currently has Distributions associated with this Channel, and all Subscribers currently subscribed to the Channel have previously received a certificate from this Distributor, select No.
- ◆ If this is the first Distribution added to this Channel by the Distributor, or a Subscriber has been newly added to the Channel, select Yes (to resolve certificates).

This copies the security certificates from the Distributor to the Subscribers subscribed to the Channel.

- ◆ If the server is a Linux or Solaris Subscriber that does not have a drive mapped to it (such as through using Samba) from the workstation you are using to resolve certificates, see [Section 7.1.10, "Manually Copying Certificates for Non-Encrypted Distributions," on page 308.](#)

A prompt to copy a certificate is usually displayed when you have added:

- ◆ A Channel to a Distribution
- ◆ A Distribution to a Channel
- ◆ A Subscriber to a Channel
- ◆ A Channel to a Subscriber

To initiate resolving certificates:

- 1 In ConsoleOne, right-click the Distributor object, then click *Resolve Certificates*.
- 2 Make sure the *Copy Certificates Automatically to Subscribers* option is selected, then click *OK*.

This copies the new certificate to each Subscriber so that it can receive Distributions from this Distributor, as long as the workstation where you are running ConsoleOne can contact all of the Subscriber servers. If you are prompted for a location to copy the certificates, you must have a drive mapped to the destination server.

For information specific to resolving certificates for External Subscribers, see [Section 7.1.10, “Manually Copying Certificates for Non-Encrypted Distributions,” on page 308](#).

7.1.7 Handling Invalid Certificates

A Subscriber cannot receive Distributions from a Distributor when the Distributor’s certificate has become invalid. A Subscriber cannot receive encrypted Distributions when the Subscriber’s encryption certificate has become invalid. For information on encryption certificates, see [Section 7.2, “Distribution Security Using Encryption,” on page 309](#).

A Distributor’s certificate can become invalid when the DNS name or IP address of the Distributor has been changed. However, if your Distributor is configured to use DNS (the recommended addressing method), IP address changes on the Distributor do not invalidate its certificate. Also, if DNS addressing is being used, changes in a Subscriber’s DNS name or IP address do not prevent the Subscriber from receiving Distributions.

However, a Subscriber’s encryption certificate can become invalid when the DNS name or IP address of the Subscriber is changed, in which case a new encryption certificate needs to be created.

The following applies for DNS name changes where DNS is your installed addressing method, or for IP address changes where IP address is your installed addressing method:

- ◆ [“Distributor DNS Name or IP Address Is Changed” on page 304](#)
- ◆ [“Subscriber DNS Name or IP Address Is Changed” on page 306](#)

Distributor DNS Name or IP Address Is Changed

Because the Distributor identifies itself to Subscribers by its server’s DNS name or IP address, if you change the identifier being used on the Distributor server, Subscribers do not recognize the Distributor as a valid source for Distributions.

Changing the DNS name or IP address of a Distributor causes the certificate created by the Distributor to be invalid for all Subscribers that have received the certificate from this Distributor. Therefore, the Distributor must send new certificates to all Subscribers receiving Distributions from that Distributor.

To re-create and resolve the Distributor’s certificate, do the following in order:

1. [“Modify the Distributor Server’s Identification Attributes” on page 304](#)
2. [“Create and Send New Certificates” on page 305](#)

Modify the Distributor Server’s Identification Attributes

You must first modify the Network Address attribute on the Other tab in the Distributor and Subscriber objects’ properties.

If the server is using the DNS Name attribute to identify itself, do the following:

- 1 In ConsoleOne, right-click the Distributor object, click *Properties*, then select the *Other* tab.

- 2 Click the + symbol to the left of *NetWork Address*.
- 3 Select the icon to the left of the field you want to modify.
A *Browse* button is displayed to the right.
- 4 Click the *Browse* button.
- 5 If you are modifying the *DNS Name* field, click the drop-down list at the top of the box where Type 13 is displayed.
- 6 Change the value from Type 13 to IP, then change IP back to Type 13.
This resets the value to now recognize the new DNS name.
- 7 Click the *Browse* button to the right of the *NetAddress* field in the lower portion of the box.
- 8 Select *Servers DNS Name* (on the right side of the box), then change it to the new name.
- 9 Click *OK* to return to the *Other* tab.
- 10 Click *OK* to finish.

If the server is using the IP Address attribute to identify itself, do the following:

- 1 In ConsoleOne, right-click the Distributor object, click *Properties*, then select the *Other* tab.
- 2 Click the + symbol to the left of *NetWork Address*.
- 3 Select the icon to the left of the field you want to modify.
A *Browse* button is displayed to the right.
- 4 Click the *Browse* button.
The IP address is displayed in the lower portion of the dialog box.
- 5 Change the IP address to the new one.
- 6 Click *OK* to return to the *Other* tab.
- 7 Click *OK* to finish.

Continue with [“Create and Send New Certificates” on page 305](#).

Create and Send New Certificates

- 1 On the Distributor server, shut down the Distributor Agent:
NetWare: At the ZENworks Server Management console prompt, enter `exit`.
Windows: In the Services dialog box, stop the Novell ZENworks Service Manager service.
For information on stopping and starting agents, see [“Starting and Stopping Server Management Services”](#) in the *Novell ZENworks 7 Server Management Installation Guide*.
- 2 In the `\zenworks\pds\ted\security\private` directory on the Distributor server, delete the `.keystore` file.
This file contains the Distributor’s certificate.
- 3 In the `\zenworks\pds\ted\security\csr` directory on the Distributor server, delete the `.csr` file that has a name that matches either the old DNS name or the old IP address.
- 4 Restart the Distributor Agent.
A new certificate and `.keystore` file are automatically created for the Distributor.

5 To send new certificates to all Subscribers that receive Distributions from the Distributor selected in [Step 1](#):

5a To resolve certificates, in ConsoleOne, right-click the Distributor object, then click Resolve Certificates.

IMPORTANT: ConsoleOne copies the certificate files to Subscriber servers. Therefore, the client software on the workstation running ConsoleOne must have access to the Subscriber servers' file systems. For Windows Subscriber servers, the Domain and Workgroup rights on the workstation must be set up to facilitate automatic certificate copying. Otherwise, a 1204a error is given.

5b Make sure the Copy Certificates Automatically to Subscribers option is selected, then click OK.

This copies the new certificate to each Subscriber so that it can receive Distributions from this Distributor, as long as the workstation where you are running ConsoleOne can contact all of the Subscriber servers. If you are prompted for a location to copy the certificates, you must have a drive mapped to the destination server.

Subscriber DNS Name or IP Address Is Changed

Because the Distributor obtains the address of a Subscribers from the Subscriber's object in eDirectory, this information must be updated in the Subscriber object so that it can receive its Distributions.

Changing the DNS name or IP address of a Subscriber causes all encryption certificates contained on the Subscriber to be invalid. Subscribers can have one encryption certificate from each Distributor that sends it encrypted Distributions.

Subscribers can continue to receive non-encrypted Distributions, even if the DNS name or IP address is changed.

The following sections outline the steps to resolve DNS name or IP address changes:

- ◆ [“Modify the Subscriber Server's Identification Attributes” on page 306](#)
- ◆ [“Resolve the New Certificates” on page 307](#)

Modify the Subscriber Server's Identification Attributes

You must first modify the Network Address attribute on the Other page in the Distributor and Subscriber objects' properties. To accomplish this, do the following as applicable.

If the server is using the DNS Name attribute to identify itself, do the following:

- 1** In ConsoleOne, right-click the Subscriber object, click *Properties*, then select the *Other* tab.
- 2** Click the + symbol to the left of *NetWork Address*.
- 3** Select the icon to the left of the field you want to modify.
A *Browse* button is displayed to the right.
- 4** Click the *Browse* button.
- 5** If you are modifying the *DNS Name* field, click the drop-down list at the top of the box where Type 13 is displayed.
- 6** Change the value from Type 13 to IP, then change IP back to Type 13.

This resets the value to now recognize the new DNS name.

- 7 Click the *Browse* button to the right of the *NetAddress* field in the lower portion of the box.
- 8 Click *Servers DNS Name* (on the right side of the box), then change it to the new name.
- 9 Click *OK* to return to the *Other* tab.
- 10 Click *OK* to finish.

If the server is using the IP Address attribute to identify itself, do the following:

- 1 In ConsoleOne, right-click the Subscriber object, click *Properties*, then select the *Other* tab.
- 2 Click the + symbol to the left of *NetWork Address*.
- 3 Select the icon to the left of the field you want to modify.
A *Browse* button is displayed to the right.
- 4 Click the *Browse* button.
The IP address is displayed in the lower portion of the dialog box.
- 5 Change the IP address to the new one.
- 6 Click *OK* to return to the *Other* tab.
- 7 Click *OK* to finish.

Resolve the New Certificates

To reproduce valid encryption certificates for the Subscriber, follow the instructions under [Section 7.2, “Distribution Security Using Encryption,” on page 309](#).

7.1.8 Certificate and Private Key Directories

Certificates and private keys for Policy and Distribution Services are stored in the following locations in the `.keystore` file:

- ◆ For the Distributor’s private key on a NetWare Distributor server:

```
sys:\zenworks\pds\ted\security\private
```

- ◆ For the Distributor’s private key on a Windows Subscriber server:

```
c:\zenworks\pds\ted\security\private
```

- ◆ For certificates received from Distributors on a NetWare Subscriber server:

```
sys:\zenworks\pds\ted\security
```

After the Distribution has been sent, the certificate is moved into the `.keystore` file.

7.1.9 Creating Security Certificates for Non-Encrypted Distributions

To create a certificate on a Distributor and copy it to its associated Subscribers:

- 1 On the server where a Distributor is installed, make sure its Distributor Agent is running (use `zfs.ncf` on a NetWare server, restart the Novell ZENworks Service Manager service on a Windows server, or enter `/etc/init.d/novell-zfs start` on a Linux or Solaris server).

This Java process creates the certificate and writes it to eDirectory.

- 2 Copy the certificate to each Subscriber using one of the following methods:
 - ♦ If your Channels and Distributions are set up, right-click the Distributor object in ConsoleOne, click *Resolve Certificates*, then click *OK*. Make sure the *Copy Certificates Automatically to Subscribers* option is selected before clicking *OK*. This copies the new certificate to each Subscriber so that it can receive Distributions from this Distributor.
For information on resolving certificates, see [Section 7.1.6, “Resolving Certificates,” on page 303](#).
 - ♦ If necessary, associate Subscribers with a Channel, create a Distribution for the Distributor, then associate the Distribution with a Channel. When you click *OK*, you are prompted to resolve the certificate. Respond to the query with *Yes* to resolve certificates for all Subscribers. The certificates are copied to all of the associated Subscribers. The Subscriber Java process does not need to be running on the Subscriber server; the server only needs to be up.
 - ♦ Manually copy the Distributor’s certificate to each Subscriber server’s `installation_path\zenworks\pds\ted\security` directory (on Linux or Solaris, `/var/opt/novell/zenworks/zfs/pds/ted/security`). This method is necessary if you do not have a drive mapped to the Linux or Solaris server to the workstation you are using to resolve certificates.
 - ♦ Right-click a Subscriber object, then click *Resolve Certificates* (repeat for each Subscriber object). This option might only be available if you answered *No* when prompted to copy security certificates.

The first two options are the easiest when there are many Subscribers receiving Distributions from one Distributor.

- 3 Because each Distributor creates its own security certificate, repeat [Step 1](#) and [Step 2](#) for each Distributor object in the tree.

7.1.10 Manually Copying Certificates for Non-Encrypted Distributions

To manually copy certificates to Subscribers using ConsoleOne:

- 1 Right-click a Distributor, Subscriber, or External Subscriber object, then click *Resolve Certificates*.

or

Click *File*, then click *Resolve Certificates*.

- 2 Select the *Save Certificates to Disk* option.
- 3 Provide a path for where to copy the certificate file, then click *OK*.

The certificate file that is copied to this path is named using the following syntax:

`DNS_Name.cer`

- 4 Copy the `DNS_name.cer` file from the path you gave to the Subscriber server’s `\zenworks\pds\ted\security` directory (on Linux or Solaris, `/var/opt/novell/zenworks/zfs/pds/ted/security`).

7.2 Distribution Security Using Encryption

Policy and Distribution Services provides the option to encrypt a Distribution to prevent unauthorized access to its contents when the Distribution is sent outside your secured network. There is usually no need to encrypt Distributions that are sent within your secured network.

Encrypting Distributions is a two-step process:

1. Select the Encrypt check box in the Distribution's properties in ConsoleOne and select the level of encryption (strong or weak).
2. Manually create and copy the encryption security certificate files between the Distributor and Subscriber servers.

IMPORTANT: For security, you should use a physical medium, such as a diskette, to transfer the certificate between network servers.

Thereafter, the Distribution is sent as an encrypted Distribution.

To understand Distribution encryption, review the following:

- ◆ [Section 7.2.1, "Creating and Copying Encryption Certificates," on page 309](#)
- ◆ [Section 7.2.2, "Sending an Encrypted Distribution," on page 311](#)
- ◆ [Section 7.2.3, "Extracting an Encrypted Distribution," on page 312](#)

7.2.1 Creating and Copying Encryption Certificates

RSA PKIs provide the security process used for encrypted Distributions.

Encryption certificates are created from Certificate Signing Request (.csr) files. Every Subscriber server contains a .csr file that can be used as a template for creating an encryption certificate for a particular Distributor.

The encryption certificates (.cer) are used by the Subscribers to ensure secure transmission of an encrypted Distribution. If you pass the .cer file over the wire, the Distribution's encryption key could be compromised. Therefore, you must manually copy the encryption security certificates to ensure that the encryption key contained in the certificate files is kept secure.

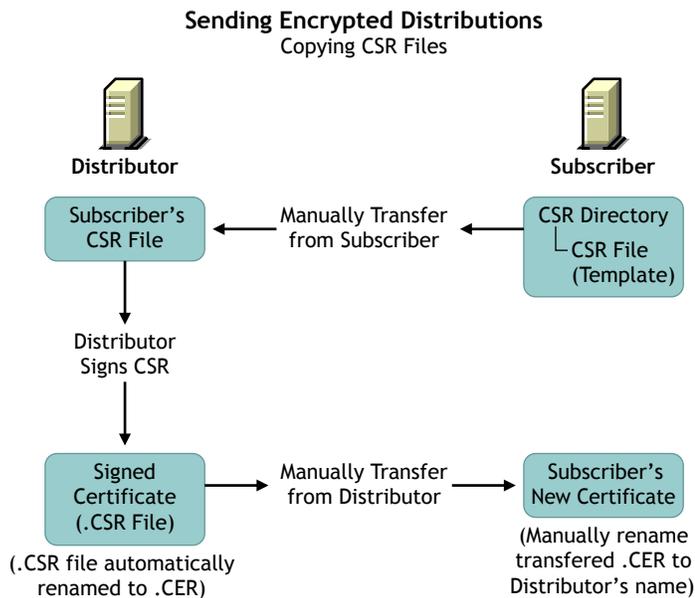
IMPORTANT: Do not manually copy a certificate by using a file browser, because that uses transmission lines and can be compromised. Instead, copy the certificate to an external media, such as a floppy diskette, and transport it physically between the Distributor and Subscriber servers.

To use encryption certificates with Subscribers, you must have previously resolved certificates and sent a non-encrypted Distribution to each Subscriber.

For information on resolving certificates, see [Section 7.1.6, "Resolving Certificates," on page 303](#).

[Figure 7-3](#) illustrates the process of manually copying the encryption certificates:

Figure 7-3 Manually Copying Encryption Certificates



The Distributor signs the `.csr` to create the encryption `.cer` file, which is manually copied from the Distributor to the Subscriber to replace the current non-encryption `.cer` file on the Subscriber server.

The encryption certificate is required for extracting a Distribution. If a Subscriber is only acting as a parent Subscriber to pass the encrypted Distribution on to Subscribers who have subscribed to the Distribution's Channel, the parent Subscriber does not need to have the encryption certificate on its server.

To create certificates for an encrypted Distribution:

- 1 Determine the Distribution you want encrypted.
- 2 Determine the Distributor that owns this Distribution.
- 3 Determine which Subscribers should receive the encrypted Distribution.
- 4 Resolve certificates for the selected Distributor to the selected Subscribers, then send a non-encrypted Distribution from that Distributor to the Subscribers.

For information on resolving certificates, see [Section 7.1.6, "Resolving Certificates," on page 303](#).

- 5 Access the file systems of this Distributor and these Subscribers.
- 6 Copy every `.csr` certificate file contained in the following directory from each Subscriber to the same path on the Distributor:

```
\zenworks\pds\ted\security\csr
```

This path begins with whatever you used for installing ZENworks Server Management.

The Certificate Signing Request (`.csr`) is used to create the encryption certificate file.

- 7 In ConsoleOne, right-click the Distributor object, click *Sign CSR Files*, select the `.csr` files to be signed, click *Sign*, click *OK* on the Success dialog box, then click *Close*.

You can select multiple `.csr` files to be signed at the same time.

This creates the Certificate (.cer) files in the same Distributor's directory as the .csr files you copied from the Subscribers. You will have one .cer file for each .csr file.

You can also perform this step using iManager:

- 7a** Select *Remote Web Console*.
 - 7b** Select or provide the Distributor's IP address.
 - 7c** In the *Available Services* drop-down box, select *Tiered Electronic Distribution*.
 - 7d** Select the *Security* tab, then click the *Sign CSR* link.
- 8** For each target Subscriber, do the following:
- 8a** Copy the Subscriber server's corresponding .cer files from the following location on the Distributor's file system:
`\zenworks\pds\ted\security\csr`
to the following path on the Subscriber's own server's file system:
`\zenworks\pds\ted\security`
Each .cer file contains its Subscriber server's name.
 - 8b** Rename the .cer files that you just copied to the Subscriber server to have the Distributor's DNS name instead of the Subscriber's.
- 9** Send the encrypted Distribution.

WARNING: Under the following scenario, the encryption certificates you just created can be overwritten before they are used:

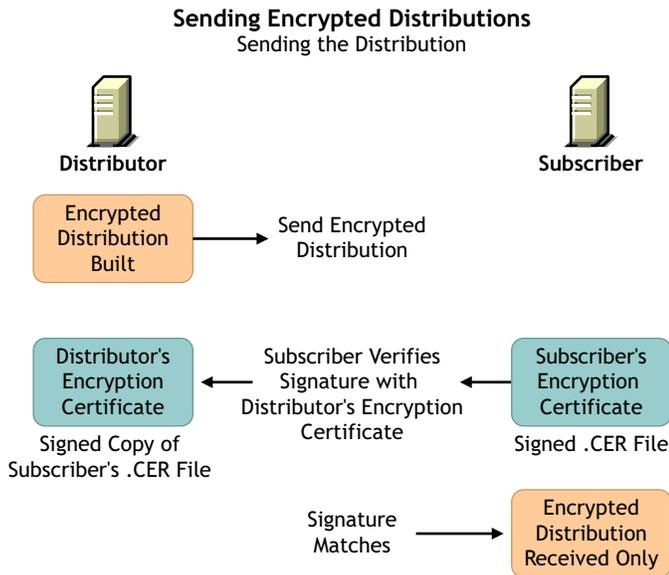
1. Changes are made to the Channel, Subscribers, or Distribution involved with the encrypted Distribution.
2. This causes the prompt for copying certificates to be displayed.
3. If you reply with Yes before the encrypted Distribution has been sent and received by the Subscribers:
 - a. The encryption .cer file is overwritten on each Subscriber with a non-encryption .cer file.
 - b. The Subscribers cannot decrypt the Distribution when it is received, because the .cer file was overwritten with a .cer file that does not contain the encryption keys.

After the encrypted Distribution has been sent once to each Subscriber, the encryption .cer file is moved into the .keystore file on the Subscriber server's file system so that it cannot be overwritten. Thereafter, you can reply with Yes to copy certificates when this scenario occurs.

7.2.2 Sending an Encrypted Distribution

After an encryption certificate has been established on a Subscriber server, [Figure 7-4](#) illustrates the process for sending encrypted Distributions:

Figure 7-4 Sending Encrypted Distributions

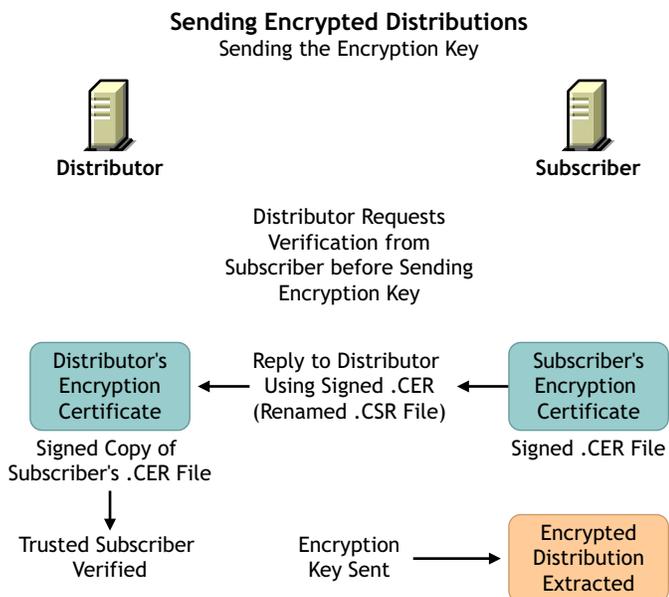


The only Subscribers that need to receive the encryption key are those that are extracting the Distribution. Therefore, parent Subscribers and Subscribers in the Distributor's routing hierarchy do not need to receive the encryption key if they are not extracting the Distribution.

7.2.3 Extracting an Encrypted Distribution

Before an encrypted Distribution can be extracted on a Subscriber server, the Subscriber must receive the encryption key. Figure 7-5 illustrates how the key is sent:

Figure 7-5 Sending Encryption Keys



Each Distribution has its own encryption key sent.

7.3 Security for Inter-Server Communication Across Non-Secured Connections

Policy and Distribution Services uses XMLRPC (Extensible Markup Language Remote Procedure Call) for its normal inter-server communications. XMLRPC optionally provides security for inter-server communication across non-secured connections. Policy and Distribution Services can use this security for inter-server communications between servers across non-secured connections, or between a management workstation and servers across non-secured connections. For example, firewalls, intranets, or NAT configurations.

This inter-server communications security ensures that data received across a non-secured connection is from a trusted source, that it has not been tampered with en route, and that the data received can be trusted by other machines. This is accomplished through the use of signed security certificates and digital signatures.

This security requires modifications to certain text files, and is installed using a Server Management wizard.

The following are instances when you could want inter-server communication security:

- ♦ **ConsoleOne administration:** When you use a workstation to manage a Distributor server across a non-secured connection.
- ♦ **SET parameters:** When you create a SET Parameter policy or a software package for SET parameters, inter-server communication takes place to provide the target server's SET parameter information. This communication could cross a non-secured connection.
- ♦ **Server Down policy:** When you use this policy to down a server, the communication between the downed server and another server watching for it to come back up could cross a non-secured connection.

For instructions on installing XMLRPC security, see “[Installing Additional Security for Non-Secured Connections](#)” in the *Novell ZENworks 7 Server Management Installation Guide*.

Review the following sections to understand inter-server communications security using XMLRPC:

- ♦ [Section 7.3.1, “Terms Used in This Section,” on page 313](#)
- ♦ [Section 7.3.2, “Security Certificates,” on page 314](#)
- ♦ [Section 7.3.3, “Using SSL,” on page 314](#)
- ♦ [Section 7.3.4, “Format of the Password File,” on page 314](#)
- ♦ [Section 7.3.5, “TCP/IP Addresses and DNS Names,” on page 315](#)

7.3.1 Terms Used in This Section

The terms and acronyms listed in [Table 7-1](#) are used in this security documentation:

Table 7-1 *Inter-Server Communications Security Terms*

| Term | Explanation |
|---|--|
| CA | Certificate Authority The trusted certificate source responsible for digitally signing other server's x.509 certificates. |
| CS | Certificate Signer The trusted certificate source responsible for digitally signing other server's XMLRPC certificates. |
| certificate or security certificate | An electronic document that contains an electronic signature for validating anything associated with the certificate, such as a Distribution. |
| CSR | Certificate Signing Request Request by a server to have an XMLRPC certificate signed by the trusted CS. This is not an X.509 certificate that would be signed by a root CA, such as VeriSign* or Thawte Consulting. |
| self-signed certificate | A valid certificate signed by its creator. |
| signed certificate | A certificate signed by a CS, which makes it valid for acceptance by the receiving server. |
| SSL | Secure Socket Layer |
| XMLRPC | Extensible Markup Language Remote Procedure Call Software used by Server Management and Tiered Electronic Distribution for inter-server communications. |

7.3.2 Security Certificates

Inter-server communications security uses signed certificates issued by the Certificate Signer (CS), which are valid only within the context of the Novell ZENworks family of products.

The certificates used are not X.509 compliant and cannot be used for any e-commerce or SSL applications.

7.3.3 Using SSL

When a CS servlet signs a Certificate Signing Request (CSR), the requesting client must authenticate with a username and password via HTTP Basic Authentication. You can secure the username and password by using SSL. For information on how to enable SSL for a commercial Web server, see your SSL documentation.

7.3.4 Format of the Password File

Inter-server communications security uses a password file for the username and password that are authenticated for CSR signing. You can create the password file in a text editor and place it in any secure location. You should also restrict access to the file to only the users who are listed in the file.

Username and passwords are both case sensitive. The syntax for the password file is:

```
username=password
```

For example:

```
admin=adminpassword
```

```
CSsigner=cspassword
```

```
JohnDoe=jdpassword
```

You should limit the access to the password file to those users included within the file.

7.3.5 TCP/IP Addresses and DNS Names

In setting up inter-server communications security, the installation program relies on addresses or names of the servers where you want this security enabled. You can use either TCP/IP addresses or fully distinguished DNS server names.

For the various methods you can use to obtain these addresses or server names, see “[Gather Information for Installation](#)” in the *Novell ZENworks 7 Server Management Installation Guide*.

The following information on scheduling applies to Policy and Distribution Services in Novell® ZENworks® Server Management:

- ♦ [Section 8.1, “Understanding Scheduling in Policy and Distribution Services,” on page 317](#)
- ♦ [Section 8.2, “Scheduling and Tiered Electronic Distribution Objects,” on page 319](#)
- ♦ [Section 8.3, “Scheduling and Server Policies,” on page 338](#)

8.1 Understanding Scheduling in Policy and Distribution Services

Review the following:

- ♦ [Section 8.1.1, “Why Scheduling is Necessary for Distributions,” on page 317](#)
- ♦ [Section 8.1.2, “Scheduling Is Required for Some Server Policies,” on page 318](#)
- ♦ [Section 8.1.3, “Scheduling Differences Between Server Policies and Tiered Electronic Distribution,” on page 318](#)
- ♦ [Section 8.1.4, “Precedence of the Tiered Electronic Distribution Policy,” on page 319](#)

8.1.1 Why Scheduling is Necessary for Distributions

When you create a Distribution (by creating and configuring its object), you want it to reach certain Subscriber servers to be used by them, and you want that to happen in a useful time frame. The distribution process requires scheduling in order to do this.

Basically, the distribution process is:

1. You create and configure a Distribution object.
2. The Distributor object that you assigned the Distribution object to reads eDirectory and discovers the new Distribution.
3. The Distributor server builds the Distribution file according to your configuration.
4. You associate the Distribution object with a Channel object.
5. The Distributor server sends the Distribution file to the Subscriber servers that are subscribed to that Channel.
6. The Subscriber servers extract and install the Distribution’s content.

This distribution sequence needs to be scheduled for the following reasons:

- ♦ **Ordering the distribution process:** The flow of a Distribution from one process to another needs to be ordered so that the Distribution gets distributed and used in a timely manner.

Conflicting scheduling might cause a Distribution to never get through the process, or to arrive and get used much later than you anticipated.

- ♦ **Minimizing network traffic:** Scheduling can provide flexibility in controlling network bandwidth usage. For example, you can schedule large Distributions to be sent when your network's traffic is at its lightest.
- ♦ **Minimizing impact on servers:** Scheduling helps to minimize the impact of building, sending, and extracting Distributions for the servers involved. For example, you can schedule large Distributions to be built and extracted during off-peak hours or on weekends.

Scheduling does not affect the total network resources used by a Distribution. It only affects when those resources are used.

8.1.2 Scheduling Is Required for Some Server Policies

Some policies must be scheduled before they can be enforced.

If you enable a policy, but do not schedule it, it is activated according to the schedule currently specified in the Default Package Schedule, which provides a default for scheduled policies. The default schedule is to run at System Startup.

The order of enforcement of different server policies is not guaranteed if the policies use exactly the same schedule. In other words, you should stagger the policies' schedules if you want to ensure the order in which they are enforced.

For information on scheduling policies, see [Section 8.3, "Scheduling and Server Policies," on page 338](#).

For information on policies, see [Chapter 4, "Server Policies," on page 193](#).

8.1.3 Scheduling Differences Between Server Policies and Tiered Electronic Distribution

Policies are scheduled according to local times. Tiered Electronic Distribution objects are scheduled according to an offset from Greenwich Mean Time (GMT).

Server Policies example: If you are residing in Utah and set a policy to be executed at 5 p.m. Utah time, it would be executed at 5 p.m. local time in Utah for servers residing in Utah. In California, it would execute at 5 p.m. local time in California. In other words, setting a time of 5 p.m. for a policy makes it execute at 5 p.m. local time wherever the servers reside.

Tiered Electronic Distribution example: If you are residing in Utah during Daylight Saving Time and set a Tiered Electronic Distribution object's schedule for 5 p.m., it would be executed at 5 p.m. local time in Utah. In California, it would execute at 4 p.m. local time (5 p.m. in Utah) for servers residing in California. In other words, Tiered Electronic Distribution schedules are relative to a GMT offset that makes the Tiered Electronic Distribution schedule execute at the exact same moment worldwide.

For Distributions, you can define a window of opportunity during the day for when a schedule's action is to begin and end. Distributions are anticipated to occur during off-peak hours. For some networks, it is possible that the scheduling window can be very short. Other systems on the network also use off-peak hours for processing, such as backups.

You can have instances where the limiting factor is available time; therefore, the critical condition is how fast the distributions can take place, regardless of the resources consumed. You might need to experiment to determine the best relationship between time and resources.

8.1.4 Precedence of the Tiered Electronic Distribution Policy

If you set a schedule in the Schedule tab for the Tiered Electronic Distribution policy (in the Service Location Package), this schedule is the default for all Distributors and Subscribers for which the policy applies, unless in ConsoleOne you set a schedule for a specific Tiered Electronic Distribution object. In other words, modified schedules for Distributors and Subscribers automatically override the Tiered Electronic Distribution policy schedule.

The Distributor and Subscriber schedules are different. There are separate Schedule tabs for the Distributor's Refresh and Subscriber's Extract schedules.

By default, when a schedule is set in the Tiered Electronic Distribution policy, the Use Policy check boxes are displayed on both the General and Schedule tabs for all Distributors and Subscribers. And, the box is automatically selected for the Distributor and Subscriber objects that have not yet had their schedules modified. It is deselected for the objects that have a schedule defined.

You can disable the Tiered Electronic Distribution policy's default schedule for a specific Distributor or Subscriber by deselecting the Use Policy check box in the object's properties. Then you must define a schedule in the object's properties for it to have a usable schedule.

You can override a specific Distributor or Subscriber schedule by selecting the Use Policy check box in that object's properties. The Tiered Electronic Distribution policy's schedule is then applied to that Distributor or Subscriber.

For information on how to create, configure, and schedule the Tiered Electronic Distribution policy, see [“Tiered Electronic Distribution” on page 209](#).

8.2 Scheduling and Tiered Electronic Distribution Objects

Scheduling can be a complex undertaking if you do not understand the fundamental scheduling principles. Review the following sections for guidelines that will help you to set up effective schedules for your Distributions:

- ◆ [Section 8.2.1, “Understanding the Tiered Electronic Distribution Objects and Their Schedules,” on page 319](#)
- ◆ [Section 8.2.2, “How the Tiered Electronic Distribution Schedules Interrelate,” on page 324](#)
- ◆ [Section 8.2.3, “The Three Timing Aspects of Scheduling,” on page 326](#)
- ◆ [Section 8.2.4, “Approaches to Scheduling,” on page 329](#)
- ◆ [Section 8.2.5, “Scheduling Issues,” on page 331](#)

8.2.1 Understanding the Tiered Electronic Distribution Objects and Their Schedules

When sending Distributions between Distributor and Subscriber servers, several Tiered Electronic Distribution objects are involved in the distribution process. Because of this, you must set schedules in some of the objects so that the process flows efficiently, yielding the intended distribution results.

The following sections explain the schedules:

- ◆ [“The Tiered Electronic Distribution Schedules” on page 320](#)

- ◆ “Distributor Object’s Refresh Schedule” on page 321
- ◆ “Distribution Object’s Build Schedule” on page 322
- ◆ “Channel Object’s Send Schedule” on page 323
- ◆ “Subscriber Object’s Extract Schedule” on page 324

The Tiered Electronic Distribution Schedules

The Tiered Electronic Distribution objects listed in [Table 8-1](#) can be scheduled:

Table 8-1 Tiered Electronic Distribution Schedules

| Tiered Electronic Distribution Object | Schedule Name | Scheduling Purpose |
|---------------------------------------|---------------|--|
| Distributor | Refresh | Tells the Distributor when it should re-read eDirectory to discover any changes to its Distributions. If it finds changes, it rebuilds the Distributions according to the Distribution objects’ Build schedules. |
| Distribution | Build | Tells the Distributor when it can build a particular Distribution. |
| Channel | Send | Tells the Distributor when it can send the Distributions it owns in the Channel. |
| Subscriber | Extract | Tells the Subscriber when it can extract and install any Distributions it has received and hasn’t yet extracted. |

The above Tiered Electronic Distribution objects must be scheduled or they cannot perform their Distribution-related actions, such as determining when Distributions are discovered, built, distributed, and extracted.

The following Tiered Electronic Distribution objects do not have schedules:

- External Subscriber
- Subscriber Group
- Policy Package ¹

¹ Only the Container Package and Service Location Package. The Distributed Policy Package can be scheduled using the Schedule tab in the Distribution object.

The following sections explain scheduling issues:

- ◆ “Server CPU Usage by the Schedules” on page 320
- ◆ “Schedule Types” on page 321
- ◆ “Resolving Certificates when Changing Schedules” on page 321

Server CPU Usage by the Schedules

Schedules do not directly affect the total resources used by a Distribution (such as CPU cycles, bandwidth, and disk space), but rather when the resources are used. Therefore, Tiered Electronic Distribution’s schedules control when Distributions are built, sent, and extracted.

However, CPU usage is affected by which servers are being used to perform a schedule's action. A server's CPU time depends on which Tiered Electronic Distribution function is running on the server, as shown in [Table 8-2](#):

Table 8-2 CPU Time Usage by Schedule

| Server's CPU Time | Schedules |
|-------------------|----------------------|
| Distributor | Refresh, Build, Send |
| Subscriber | Extract |
| parent Subscriber | Send, Extract |

Schedule Types

When you set an object's schedule, you have the following schedule types to choose from:

- Never
- Daily
- Monthly
- Yearly
- Interval
- Time
- Run Immediately (except for the Distributor object)

For more information on the schedule types, see [“Frequency” on page 326](#) and [Appendix B, “Schedule Types,” on page 399](#).

Resolving Certificates when Changing Schedules

You might need to resolve certificates when making changes to one of the schedules. For more information, see [Section 7.1.6, “Resolving Certificates,” on page 303](#).

Distributor Object's Refresh Schedule

A Distributor's schedule determines when the Distributor reads Novell eDirectory™ for configuration changes. This enables the Distributor to respond to a request to build a Distribution. The Distributor rebuilds a Distribution when the Distribution's schedule indicates that it should be built.

When the Channel's Send schedule starts, the Distributor checks with the Subscribers that it sends to directly to see if they have the current Distribution. However:

- ◆ If the Distribution is non-sequential, the Distributor simply checks for the current version.
- ◆ If the Distribution is sequential (the File or Desktop Application types of Distributions only), it checks to see if the Subscribers have all of the versions of the Distribution, starting with the baseline and every change since the baseline.

If the Subscriber does have the entire Distribution, it checks with its subordinate Subscribers to see if they do, and so on down the routing hierarchy.

The time it takes to verify that all receivers have all of the Distributions in the Channel is minimal.

IMPORTANT: A Distribution might never get sent completely if the Refresh schedule is shorter than the time it takes to build or send the Distribution. In other words, if the Refresh schedule is too short, when the Distributor is refreshed the Distribution in the process of being built or sent could be cancelled before it has completed sending. Therefore, we recommend the Distributor's Refresh schedule be daily, unless changes to Distributions warrant a more frequent refresh, then set it in hours. Do not refresh the Distributor more often than every five minutes.

Scheduling a Distributor

- 1 In ConsoleOne, right-click the Distributor object > click *Properties*.
- 2 Select the *Schedule* tab > click the arrow for the drop-down box > click *Interval* > select an interval, such as *Daily*.
- 3 Set the start and end times, if necessary.

The *Start Time* and the *End Time* fields specify the time window for performing the schedule's action.

You can repeat the action every so often throughout the day.

You can also have the refresh occur randomly in the specified time window. For more information, see [“Using the Randomly Dispatch Option in a Distributor's Refresh Schedule” on page 336](#).

- 4 Click *OK*.

Distribution Object's Build Schedule

The Distribution's schedule determines when a Distributor is requested to create the Distribution file based on the definition in the Distribution object.

Most Distributions consist of a set of files that change over time and need to be redistributed on a regular basis. Each Distribution has its own Build schedule that tells the Distributor how often to rebuild the Distribution. When the Distributor builds a Distribution, it automatically compares it with the previous version to see if there are any changes.

For the File Distribution, if there are no changes in the current build, no new version is created. If there are changes, a delta is built consisting of only the changes to be distributed.

For the FTP, HTTP, and Software Package Distribution types, a new version is only built if there has been a change since the last version. The Distributor sends the complete new version to all target Subscribers.

The Distribution's End Time is used to determine the end time for randomly dispatching events. In other words, the Distributor does not stop building the Distribution until it is complete.

Deleted files and directory synchronization are handled in the Build schedule.

Scheduling a Distribution

- 1 In ConsoleOne, right-click a Distribution object > click *Properties*.
- 2 Select the *Schedule* tab > click the arrow for the drop-down box > select a schedule type, such as *Run Immediately*.

You can repeat the action every so often.

The *Start Time* and the *End Time* fields specify the time window for performing the schedule's action.

You can also have the build occur randomly in the specified time window (if you select the Daily schedule type). For more information, see [“Using the Randomly Dispatch Option in a Distribution's Build Schedule” on page 337](#).

3 Click *OK*.

Channel Object's Send Schedule

A Channel's Send schedule provides a window of time for when a Distributor can start sending its Distributions to the Subscribers associated with that Channel.

The Channel's schedule applies only to the Distributor and its direct receivers (first tier Subscribers). When the Send schedule ends, the Distributor stops distributing to those first tier Subscribers.

Second-tier receivers and beyond do not adhere to the Channel's schedule. The parent Subscribers that are sending Distributions to other Subscribers continues to send a Distribution after the Send schedule ends. Their subordinate Subscribers also ignore the Send schedule.

The Send schedule's End Time forces the Distributor to stop sending a Distribution when the Send schedule ends. The Distributor starts sending the Distribution where it left off when the Send schedule begins again. A Distribution is not totally re-sent. For example, if 50 MB of a 60 MB Distribution had already been sent before the disruption, when the Send schedule starts again for the Channel, the Distributor begins sending the remaining 10 MBs.

For information on how time zones affect a Channel's schedule, see [“Scheduling Tiered Electronic Distribution Objects in Different Time Zones” on page 334](#).

Cache and Forward has no bearing on whether a parent Subscriber continues to send a Distribution when the Channel's Send schedule ends. Parent Subscribers who have completely received a Distribution prior to the Send schedule ending continues to send that Distribution to subordinate Subscribers. There is no mechanism for controlling whether parent Subscribers should continue to send when the Send schedule ends.

IMPORTANT: A Distribution might never get sent if the Send schedule is shorter than the time it takes to send the Distribution. Therefore, we recommend the Channel's Send schedule be daily or in hours. Make the Send schedule at least long enough to allow all of the Channel's Distributions to be sent.

Scheduling a Channel

- 1** In ConsoleOne, right-click the Channel object > click *Properties*.
- 2** Select the *Schedule* tab > click the arrow for the drop-down box > click *Interval* > select an interval (in the *Repeat the Action Every* field), such as 1 hour > click *OK*.

The *Start Time* and the *End Time* fields specify the time window for performing the schedule's action.

For information about randomly starting the Send schedule (if you select the Daily schedule type), see [“Using the Randomly Dispatch Option in a Channel's Send Schedule” on page 337](#).

Subscriber Object's Extract Schedule

The Subscriber's schedule determines when a Subscriber can extract a Distribution that has been received.

The Subscriber's End Time is used to determine the end time for randomly dispatching events. In other words, the Subscriber does not stop extracting the Distribution until it has completed the extraction process.

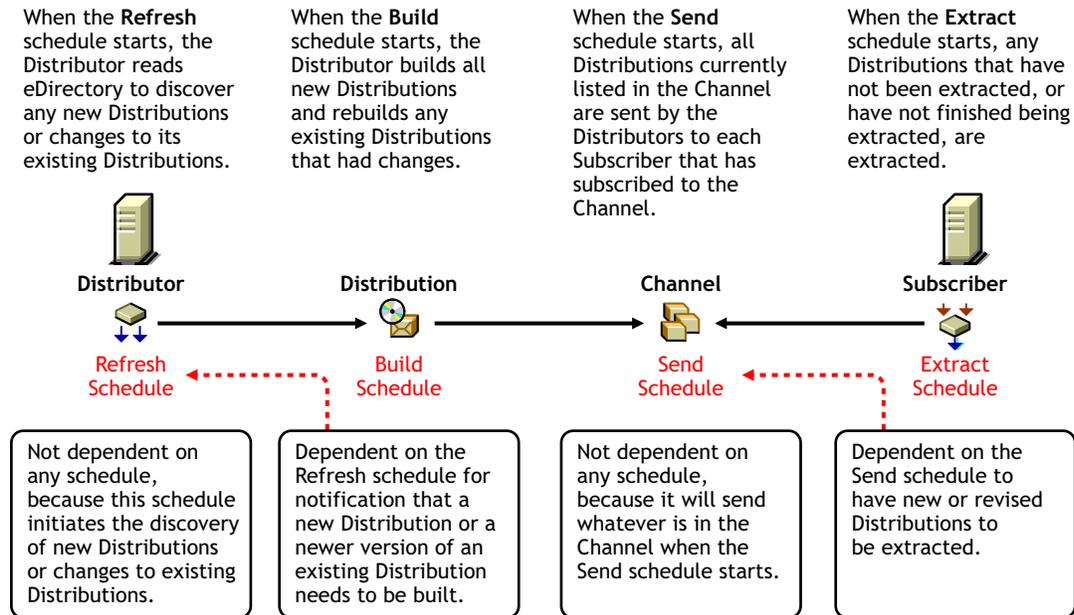
Scheduling a Subscriber

- 1 In ConsoleOne, right-click a Subscriber object > click *Properties*.
- 2 Select the *Channels* tab > click *Add* > browse for the Channel > click *Select* > click *OK*.
Make sure the Channel is listed as Active in the *Channels* list.
- 3 Select the *Schedule* tab > the arrow for the drop-down box > select a schedule, such as *Run Immediately*, then click *OK*.
This schedule type causes the Subscriber to extract the Distribution as soon as it is received.
The *Start Time* and the *End Time* fields specify the time window for performing the schedule's action.
For information about randomly starting the Extract schedule (if you select the Daily schedule type), see [“Using the Randomly Dispatch Option in a Subscriber's Extract Schedule” on page 337](#).
- 4 Repeat these steps for each Subscriber.

8.2.2 How the Tiered Electronic Distribution Schedules Interrelate

The Tiered Electronic Distribution object's schedules do not all interact directly with each other. There is a flow to how they are sequentially interconnected, as shown in [Figure 8-1](#):

Figure 8-1 Schedule Interrelationships



Most importantly, the sequence of Refresh, Build, Send, and Extract must have their schedules configured so that they allow a Distribution to be successfully discovered, built, sent, and extracted within the time frame that you intend.

If even one schedule is out of sync with the other three, it can take longer than intended for a Distribution to be created and eventually used.

For example, if you set the following schedules as indicated, the results are:

- ◆ **Set the Refresh schedule to occur hourly:** The Distributor is triggered to read eDirectory each hour to discover new or changed Distributions. You should determine this frequency by how large the Distributions are that this Distributor builds after each refresh. Generally, the shorter the time between refreshes, the better.
- ◆ **Set the Build schedule to run immediately:** The Distributor builds the new or changed Distributions immediately after the Distributor's Refresh schedule has caused it to read eDirectory to discovered them. However, very large Distributions might need to be built during off-peak hours.
- ◆ **Set the Send schedule to midnight:** The Distributor sends its Distributions during off-peak hours when the network's bandwidth is less busy.
- ◆ **Set the Send schedule to run immediately:** The Distributor sends its Distributions as soon as they are built.
- ◆ **Set the Extract schedule to 3 a.m.:** The Distributions received are extracted during off-peak hours when the server is likely to be the least busy. This is useful for servers receiving very large Distributions that do not need to be extracted and installed immediately.
- ◆ **Set the Extract schedule to run immediately:** The Subscriber extracts the Distribution as soon as it is received. This is useful for servers receiving important Distributions, such as virus patterns.

Ways that you could mess up the distribution scheduling flow:

- ◆ Setting the Distributor’s Refresh schedule to occur too frequently to allow time to build new Distributions or rebuild changed Distributions.
- ◆ Setting the Distributor’s Refresh schedule to not occur frequently enough to get important Distributions built and sent on time.
- ◆ Setting the Distribution’s Build schedule to occur too frequently to allow completion of the Distributions it was building during the previous schedule window.
- ◆ Setting the Distribution’s Build schedule to not occur frequently enough to get important Distributions built and sent on time.
- ◆ Setting the Channel’s Send schedule to not coincide with the Distributor’s Build schedule, possibly delaying the sending of Distributions.
- ◆ Setting the Channel’s Send schedule window to be too short for all its Distributions to have time to complete sending.

8.2.3 The Three Timing Aspects of Scheduling

There are three time-related aspects that affect scheduling:

- ◆ “Frequency” on page 326
- ◆ “Duration” on page 327
- ◆ “Interval” on page 328
- ◆ “Using Intervals with Distributors” on page 329

Frequency

When setting schedules, you determine how frequently you want a particular Distribution to be built, sent, and extracted.

The frequency for processing a Distribution can be determined using the schedule types listed in [Table 8-3](#):

Table 8-3 *Schedule Frequencies*

| Schedule Type | Functionality |
|-----------------|--|
| Daily | Repeats the function the same time each day |
| Monthly | Repeats the function on a specified day of the month |
| Yearly | Repeats the function on a specified day of the year |
| Interval | Repeats the function every so often (as determined by you) |
| Time | The function occurs just once at a specific date and time |
| Run Immediately | Ignores the schedule’s normal settings and starts the function immediately |

The frequency you select in scheduling the distribution process should be determined by the purpose of the Distribution. For example:

- ◆ Virus protection pattern files should be distributed and installed as soon as possible whenever they become available
- ◆ A software update should be sent and installed only once

Duration

Some schedule types have durations that you may need to determine. The duration is defined by start and end times that provide a window for the time wherein the scheduled action can be performed.

Some schedules completely stop their function at the end of the schedule's duration. Therefore, the duration of a schedule must accommodate the size of a Distribution with reference to how long it takes to build it, send it, and extract it.

Duration has different issues for different schedules, as explained in the following:

- ◆ [“The Distributor Object's Refresh Schedule” on page 327](#)
- ◆ [“The Distribution Object's Build Schedule” on page 327](#)
- ◆ [“The Channel Object's Send Schedule” on page 328](#)
- ◆ [“The Subscriber Object's Extract Schedule” on page 328](#)

The Distributor Object's Refresh Schedule

When a Refresh schedule starts, the following happens:

- ◆ **Distribution building stops:** The Distributor stops building any Distributions that it is in the middle of building. Temporary build files are not cleaned up, and building of the unfinished Distribution are not resumed where it left off. Unfinished Distributions are rebuilt by starting over when the next Build schedule starts.
- ◆ **Distribution sending is interrupted:** The Distributor stops sending any Distributions that it is in the middle of sending. However, when the Send schedule starts again, the Distributor picks up where it left off and finishes sending the Distribution.

Therefore, the Refresh schedule should not overlap the Build or Send schedules. In other words, it should start after the others end, and end when the others have not yet started.

Because the Refresh schedule can stop a Distributor from finishing a Distribution build, you may need to have multiple Distributors in your system to handle the different types of Distributions you'll be creating. For more information, see [“Determining Whether You Need Other Distributors” on page 330](#).

After the Refresh schedule's end time is reached, the Distributor picks up where it left off in sending its Distributions, but restarts building Distributions that it had not completed building when it was interrupted by the start of the Refresh schedule.

The Distribution Object's Build Schedule

When a Build schedule starts, only the Distributions that a Distributor knows about at that time start being built during the duration of the Build schedule. The Distributor learns of changes made to

existing Distributions or of newly-created Distributions by reading eDirectory, which is done according to the Distributor's Refresh schedule.

After the Build schedule's end time is reached, building continues on all Distributions that it started building until the Distributions are finished being built or failed to be built.

The Channel Object's Send Schedule

When a Send schedule starts, the Distributor begins sending its Distributions that are listed in the Channel, but only those Distributions that the Distributor knows about that are listed in the Channel at the time the Send schedule starts.

When a Send schedule's end time is reached, the Distributor stops sending its Distributions, even if the Distributions have not completed being sent. However, the next time the Send schedule starts, the Distributor picks up where it left off and completes sending the partially-sent Distributions, plus begins sending any new or revised Distributions that the Distributor discovered during its Refresh schedule time.

The Subscriber Object's Extract Schedule

When an Extract schedule starts, any Distributions it has already received or will receive during its schedule's duration begins to be extracted all at the same time.

When an Extract schedule's end time is reached, the Subscriber continues to extract all Distributions that it started to extract until the Distributions are finished being extracted or failed to be extracted.

Interval

An interval is how often during a schedule's duration that the schedule restarts its function.

An interval is the equivalent to splitting up the schedule into a consecutively run series of mini-schedules. During a schedule's duration, intervals act as a stop/start position within the duration, causing the same actions to take place as for the start and stop times of the schedule itself.

Intervals can have different issues for different schedules, as explained in the following:

- ◆ [“The Distributor Object's Refresh Schedule” on page 328](#)
- ◆ [“The Distribution Object's Build Schedule” on page 328](#)
- ◆ [“The Channel Object's Send Schedule” on page 329](#)
- ◆ [“The Subscriber Object's Extract Schedule” on page 329](#)

The Distributor Object's Refresh Schedule

Use intervals to sync up a Distributor's refresh frequency with how often you want configuration information changes Distributions, Channels, Subscribers, or policies to be recognized.

The interval should not be so short that the Distributor doesn't have time to read eDirectory and build the Distributions that it finds are new or changed.

The Distribution Object's Build Schedule

For Distributions that have changes made often to the Distribution's content that you want distributed in a timely manner, use intervals for the Build schedule to efficiently recognize those changes and provide rebuilt Distributions on time.

The Channel Object's Send Schedule

When you set intervals within a Send schedule's duration, Distributions that are in the process of being sent are stopped each time the interval begins, then pick up where it left off in sending the Distributions, plus start sending any new Distributions that were added to the queue.

If you do not use intervals, any Distributions added to the Send schedule's queue after the Send schedule starts, are not sent until the next time the Send schedule starts. Therefore, setting intervals in the Send schedule allows you to have newly-queued Distributions included in the Send schedule's window of time.

The Subscriber Object's Extract Schedule

Intervals do not make sense for the extraction process. All Distributions received prior to the start of the Extract schedule, or received while the Extract schedule is open, is extracted. Extraction continues after its schedule ends, so intervals would be ignored by the extraction process.

Using Intervals with Distributors

For any schedule type that has an interval, the event does not start until after the Distributor has re-read eDirectory. For example:

- ♦ **Daily:** If the Distributor is refreshed before the current day's time window has passed, the event runs on the current day, then every day thereafter; otherwise, it first runs during that time window on the next day, then every day thereafter.
- ♦ **Interval:** If you set the interval to be three days, the event runs three days after the day the Distributor re-reads eDirectory, then run every three days thereafter.
- ♦ **Weekly, Monthly, or Yearly:** The event runs the first day, month, or specific date (the Yearly option) after the Distributor has re-read eDirectory. For example, on Wednesday you set up a Weekly event to happen each Sunday. The Distributor re-reads eDirectory on Thursday, so the event runs the following Sunday, and every Sunday thereafter.
- ♦ **Run immediately:** As soon as the Distributor is refreshed, the event runs, then runs thereafter according to the interval you set.

To cause an event for one of the interval-related schedule types to execute out of sequence (other than Run Immediately), you can use the ZENworks Server Management role in iManager. For more information, see [Chapter 2, "Novell iManager," on page 63](#).

8.2.4 Approaches to Scheduling

The following are approaches that you can use in determining how to set up your schedules:

- ♦ ["Determining Whether You Need Other Distributors" on page 330](#)
- ♦ ["Putting Channels In Control" on page 330](#)
- ♦ ["Enabling Load-Balancing for Distributors" on page 331](#)
- ♦ ["Inactivating Schedules" on page 331](#)
- ♦ ["Scheduling Conflicts with Other Software" on page 331](#)

Determining Whether You Need Other Distributors

Distributor server workload and the ability to complete Distribution building tasks should determine how many Distributors you need.

For example, if you have a very large Distribution that you want built during off-peak hours, which does not need to be sent immediately, and also have virus pattern Distributions that do need to be sent immediately, you might need two different Distributors, one with a daily refresh schedule (because you are only going to be building the Distribution once per day), and another with a frequent refresh schedule for discovering new virus pattern changes, so that their Distributions can be built and sent on time.

Putting Channels In Control

The idea in using Tiered Electronic Distribution is that you have a Distribution that you want to be used by the Subscribers at a certain time. To do so, you would have a Distribution built at a time when you want the Subscribers to use it. The key then is to get the other schedules to cooperate in getting the new Distribution down to the Subscribers on time.

The most useful scheduling configuration to do this places emphasis on the Channel's Send schedule. Review the following scenario:

1. A Distribution's Build schedule depends on how often you expect the Distribution's information to change. For example:
 - ♦ If the Distribution consists of forms that change monthly, and it is critical to distribute the updated forms quickly, the Build schedule should be set to Daily. This means the forms would be checked each day for changes, and the change would be found the day, or within a day, of when they are made to the forms. When it is discovered that the forms have changed, a new Distribution is built.
 - ♦ If the Distribution consists of a software application that changes once every six months or so, you may want the Distribution to build weekly. When the application is changed for the Distribution, no more than a week would pass before the Distribution was rebuilt.
2. Set all of your Subscriber's Extract schedules to Run Immediately. That way, no matter when a Distribution is built and sent, the Subscriber is ready to use it.

You can have all of your Subscriber's Extract schedules set to Run Immediately and not worry about impacting the Subscriber server during peak business hours with a large Distribution, because you can use the Channel's Send schedule to control when the Subscribers receives and extracts a particular Distribution.

3. Set the Channel's Send schedule to correlate with when its Distributions are scheduled to be rebuilt, and to occur when you want the Subscribers to extract them.
 - ♦ In the case of a Distribution that changes monthly, set the Channel's Send schedule to monthly.
 - ♦ In the case of a Distribution that only changes every six months or so, set the Channel's Send schedule to yearly or at an interval of xxx number of days.
 - ♦ In the case of a large Distribution that needs to be extracted during off-peak hours, set the Channel's Send schedule to run immediately, if all you are concerned with is the Distribution's extraction, which is determined by the Subscriber's Extract schedule, which can be set to control off-peak hour extraction.

Simply, build a Distribution when it is needed, get your Subscribers ready to extract and use the Distribution as soon as they receive it, then set up your Channel to send the Distribution at the optimum time for the Subscribers.

Enabling Load-Balancing for Distributors

To help load balance a Distributor server's distribution duties, do the following:

1. Select the Maximum Number of Concurrent Distributions option on the Distribution object.
2. For the Distribution object, use the Randomly Dispatch option for the Daily, Monthly, or Yearly schedule type.

For more information on the Randomly Dispatch option, see [“Using the Randomly Dispatch During Time Period Option” on page 336](#).

This spreads the network traffic that is caused by sending many Distributions over the entire scheduling window.

Inactivating Schedules

A Distribution can be set as Active or Inactive:

- ◆ **Active:** The Active check box is found on the General tab of the Distribution object.
- ◆ **Inactive:** Inactive is used when you are building a Distribution because you want to keep it inactive until it is ready to be sent to a Subscriber.

We recommend that as you are either creating or modifying a Distribution object, its associated Channel be set to Inactive until you are ready to begin distributing the Distribution package. This prevents the Distribution from being inadvertently sent before you have completed its configuration.

Scheduling Conflicts with Other Software

Most distributions are anticipated to occur during off-peak hours. For some networks, it is possible that this scheduling window may need to be very short. Other systems on the network can also use off-peak hours for processing, such as backups.

You might have instances where the limiting factor is available time; therefore, the critical condition is how fast the Distributions can take place, regardless of the resources consumed. You might need to experiment to determine the best relationship between time and resources.

8.2.5 Scheduling Issues

The following explain various scheduling issues:

- ◆ [“Schedule Interactions” on page 332](#)
- ◆ [“Time Zones and Scheduling” on page 333](#)
- ◆ [“Using the Randomly Dispatch During Time Period Option” on page 336](#)
- ◆ [“Repeating Actions” on page 338](#)

Schedule Interactions

Each of the four schedules (Refresh, Build, Send, and Extract) interact with each other in ways that determine the success or timeliness of the distribution process:

- ◆ “Overall Interaction Issues” on page 332
- ◆ “Refresh versus Build” on page 332
- ◆ “Refresh or Build versus Send” on page 333
- ◆ “Build versus Send” on page 333
- ◆ “Send versus Extract” on page 333

Overall Interaction Issues

Because the distribution process is dependent on each of the four schedules interacting with each other in a timely manner, the purpose of a Distribution should help you to determine what each of the schedules need to be that are involved in its distribution process.

For example, if you want a virus pattern Distribution to be sent as soon as it is configured or as soon as a change to it has been completed, you need to make sure the schedules involved for all four Tiered Electronic Distribution objects allow the virus patterns to be in use by the target servers as soon as possible.

However, if you want a Distribution to be built, sent, and extracted during off-peak times, because it is very large and requires a lot of bandwidth in sending and server time in building and extracting it, then you would want each schedule to help in determining when to start those processes.

Because your Distributions can vary in both purpose and size, and because you may be using the same Distributor server for building these Distributions, you need to configure the various schedules to compensate for this. For example, the Distributor could have its Refresh schedule set to start every five minutes, and that would work for both sending Distributions immediately or during off-peak hours. This is because the Distribution’s Build schedule would trigger when the particular Distribution would get built (immediately or during off-peak hours).

Subscriber Extract schedules are where you could have conflicts between extracting Distributions immediately or during off-peak hours. However, if you set the Subscriber’s schedule to immediately extract and install its Distributions, you can use the Build and Send schedules to control when it gets certain Distributions. That way, it can extract virus pattern Distributions immediately (small, so no impact on the server), and extract large Distributions when they are sent during off-peak hours.

The fact that you can control build times individually for each Distribution, and that you usually create Channels unique to the Distributions, you can configure frequent schedules for the Distributor’s Refresh and Subscriber’s Extract schedules. In other words, most of your scheduling differences can be controlled by the Build and Send schedules.

Refresh versus Build

The Distributor builds a Distribution according to the Distribution object’s Build schedule, but not before the Distributor’s Refresh schedule has told the Distributor to read eDirectory for changes related to the Distribution, such as whether there is a new one, or that something in an existing Distribution has changed, requiring it to be rebuilt. This means the Build schedule is dependent on the Refresh schedule

Therefore, if you intend that a new Distribution be built right after you have created its object and configured it, the Distributor's Refresh schedule must be frequent enough to cause the Distribution to be built. For example, you would have the Distributor's Refresh schedule set to an interval of every five minutes.

However, if you only want a Distributor server to be building Distributions during off-peak hours, then you'd want its Refresh schedule to start and end during off-peak hours. Therefore, you would select a Refresh schedule type that allows you to specify such a time window. For example, Daily, Monthly, and Yearly each provide the capability to set a time window. Additionally, Daily allows you to specify which days of the week to read eDirectory for changes.

Refresh or Build versus Send

A Distribution can only be listed in a Channel after it has initially been built. Therefore, the Channel's Send schedule is for existing Distributions, whether they be newly built or rebuilt because of changes. The Distributor sends a Distribution according to the Channel object's Send schedule. This means the Send schedule is not dependent on the Refresh or Build schedules.

However, the Refresh and Build schedules are dependent on the Send schedule, because the Subscriber servers do not receive a Distribution until the Send schedule starts.

Build versus Send

A Channel only lists Distributions that can be sent. The Channel doesn't care whether an existing Distribution is in need of being rebuilt. The Send schedule only tells the Distributor when one or more of its Distributions can be sent to the Subscriber servers that have subscribed to the Channel.

Therefore, the Build schedule is dependent on the Send schedule in that a Distribution should be initially built or rebuilt in time to be sent when the Channel's Send schedule starts. Also, after a Distribution has been built, it must wait for the Send schedule to start to be sent.

Send versus Extract

When a Send schedule starts, the Distributor sends the Distributions listed in that Channel to the Subscriber servers subscribed to the Channel. However, the Distributions received are not extracted until the Subscriber's Extract schedule starts.

Therefore, the extraction of a Distribution is only dependent on the Extract schedule. However, the Send schedule determines when the Distribution is available for extraction. Thus, the Extract schedule is dependent on the Send schedule.

The only dependency that the Send schedule has with the Extract schedule is that a sent Distribution is not extracted and installed until the Extract schedule starts, meaning you would not count on the Send schedule alone to get Distributions completely processed.

Time Zones and Scheduling

Multiple time zones can complicate your scheduling efforts. The following sections explain some of the issues:

- ◆ [“Scheduling Tiered Electronic Distribution Objects in Different Time Zones” on page 334](#)
- ◆ [“Calculating Time Differences” on page 334](#)
- ◆ [“Using Geographically-Based Channels” on page 335](#)

Scheduling Tiered Electronic Distribution Objects in Different Time Zones

The following information concerning time zone offsets is from the perspective of the Channel object. However, this information is applicable to all Tiered Electronic Distribution objects that can be scheduled.

Because a Channel is an object in the tree that is not associated with a specific server, the Channel's time is always set to the local time zone of the workstation that is running ConsoleOne[®] and setting the Channel's schedule.

For example, if you (the administrator) live in New York City, the local time for any Channels you schedule from there is local New York time.

If Distributors in different time zones from the Channel have Distributions in that Channel, the Distributors need to send their Distributions according to the Channel's local time schedule. For example:

1. You set a Channel's schedule to be from 1 a.m. through 5 a.m. local time in Los Angeles.
2. In New York you select to have a Distributor's Distribution listed in that Los Angeles Channel.
3. The Distribution can be sent only between 4 a.m. and 8 a.m. in New York because for New York, being three hours ahead of Los Angeles, its time window of 4–8 a.m. is happening at the same time as the Los Angeles time window of 1–5 a.m.

You should use a time zone offset to determine the true local time when the Distributor can send its Distributions. Also, because a Channel's schedule determines when a Distribution can be sent, you must make sure the build schedules you set for your Distributions occur before a Channel's schedule.

Calculating Time Differences

The [World Time Server \(http://www.worldtimeserver.com\)](http://www.worldtimeserver.com) is a Web site where you can determine the time difference between any two locations in the world.

As you look at the site, note the following:

- ♦ The locations in the left frame can be listed by countries or major cities.
- ♦ The current GMT time relative to the International Date Line is displayed in the right frame.
- ♦ When you select a location in the left frame, the time displayed in the right frame includes the day, date, whether Standard Time or Daylight Saving Time is in effect, and the GMT offset.

To use this site to calculate time differences between Tiered Electronic Distribution locations,

- 1 Select the location for one of the Tiered Electronic Distribution sites.
- 2 Note the time, day/date, GMT offset, and whether Daylight Saving Time is in effect (for future reference).
- 3 Select the location for another Tiered Electronic Distribution site.
- 4 Note the time, day/date, GMT offset, and whether Daylight Saving Time is in effect.
- 5 Repeat this process for all of the Tiered Electronic Distribution locations where you want to coordinate schedules.
- 6 Using the information you have gathered, calculate the time differences between the Tiered Electronic Distribution locations.

- 7 Taking into consideration when events are taking place locally at the various Tiered Electronic Distribution locations, configure the appropriate schedules using the time differences.

As an example,

- ◆ A Distributor in Hawaii lists a Distribution in a Channel in New York.
- ◆ Using the World Time Server Web site, you can find that the offset between the two locations is -6 when Daylight Saving Time is in effect. (The negative number means it is later in the time sequence, so you must subtract Hawaii's time from New York's time to arrive at the correct a.m. or p.m.)
- ◆ If the Channel's starting time is 1 a.m. in New York, select 7 p.m. for the Distributor's schedule in Hawaii.
- ◆ The result is that the Distributor can start to send its Distribution at 7 p.m.
- ◆ Because Hawaii is not observing Daylight Saving Time and New York is, when New York moves back to Standard Time, the result would be 8 p.m.

If you wanted the Distributions to be sent later in the evening in Hawaii, the Channel's time window would have to start later than at 1 a.m. in New York. For example:

- ◆ You want the Distributions to begin sending at 11 p.m. in Hawaii.
- ◆ You need to set the Channel's start time to be 5 a.m. in New York.

When you set up your Channel schedules, you need to consider which object's time window is more important. For example, it might be more important for the Distributor to be sending Distributions during off-peak hours. Therefore, using the New York and Hawaii example, to have the Distributions begin sending after midnight Hawaii time, you would need to have the New York Channel's start time set to 6 a.m. or later.

Using Geographically-Based Channels

If you want a Distribution to be received at the same local time (such as 3 a.m.) when you have Subscribers in different time zones, use geographically-based Channels.

Using a single Channel for the Distribution only allows the Distribution to be received at the same time it is sent, meaning it could arrive at different local times if the Subscribers are not all in the same time zone. For example:

1. Distribution_A is created and listed in Channel_A.
2. The Send schedule for Channel_A is set to cause Distribution_A to be sent at 3 a.m. of the Distributor's local time.
3. Subscriber001 is in the same time zone as the Distributor.
4. Subscriber002 is in a time zone that is 2 hours later than the Distributor's.
5. Subscriber001 and Subscriber002 are subscribed to Channel_A.
6. Channel_A's Send schedule fires at 3 a.m. local time of the Distributor owning Distribution_A.
7. Distribution_A is sent at 3 a.m. local time of the Distributor to Subscriber001 and Subscriber002.
8. Subscriber001 receives the Distribution at 3 a.m. its local time.
9. Subscriber002 receives the Distribution at 1 a.m. its local time.

However, your intention is that Distribution_A be received by both Subscribers at 3 a.m. their local times.

To get Subscriber002 to also receive Distribution_A at 3 a.m. its local time instead of 1 a.m., do the following:

1. Unsubscribe Subscriber002 from Channel_A.
2. Create Channel_A2 for Distribution_A.
3. Subscribe Subscriber002 to Channel_A2.
4. Set the Send schedule for Channel_A2 to cause Distribution_A to be sent at 5 a.m. of the Distributor's local time.

Then:

1. Channel_A2's Send schedule fires at 5 a.m. local time of the Distributor owning Distribution_A.
2. Subscriber002 receives the Distribution at 3 a.m. its local time.

Using Channels that are geographically-based, the Distributor sends the same Distribution at different times of the day.

Using the Randomly Dispatch During Time Period Option

The Randomly Dispatch During Time Period option is available for each of the schedules (Distributor, Subscriber, Channel, and Distribution). It is used in conjunction with a time window (Start and End times) that you can set for a Daily, Monthly, or Yearly schedule type.

Randomly dispatching causes the scheduled action to run at any time during the window for the day. This helps load-balancing on servers. However, random-dispatched schedules can be confusing if you are expecting an action to take place immediately.

The following describe the issues for the Randomly Dispatch option:

- ◆ [“Using the Randomly Dispatch Option in a Distributor's Refresh Schedule” on page 336](#)
- ◆ [“Using the Randomly Dispatch Option in a Distribution's Build Schedule” on page 337](#)
- ◆ [“Using the Randomly Dispatch Option in a Channel's Send Schedule” on page 337](#)
- ◆ [“Using the Randomly Dispatch Option in a Subscriber's Extract Schedule” on page 337](#)

Using the Randomly Dispatch Option in a Distributor's Refresh Schedule

You can use the Randomly Dispatch option for Distributor Refresh schedules to load balance Distributor refreshes from eDirectory. This is useful to minimize the network traffic that can be caused by many Distributors trying to read eDirectory at the same time.

Be sure to coordinate a Distributor's Refresh schedule with that Distributor's related Distributions' Build and Channels' Send schedules.

The Distributor's Refresh schedule should be determined by how frequently Tiered Electronic Distribution information is updated in eDirectory. For example, how often new Distributions are created, properties of existing Distribution objects changed, new Channels are added, and so on. The Distributor cannot know of changes made to Tiered Electronic Distribution objects without re-reading eDirectory. An eDirectory refresh should finish before the Build and Send schedules begin.

IMPORTANT: Do not refresh the Distributor more often than every five minutes. The following can need up to five minutes to complete their processes: Distribution building, eDirectory replication, and tree walking (when no Search policy is defined).

If you are using the Randomly Dispatch option, you should consider the End time for the Refresh schedule when setting the Start times for the Build and Send schedules.

Using the Randomly Dispatch Option in a Distribution's Build Schedule

You can use the Randomly Dispatch option for a Distribution's Build schedule to load-balance the Distributor's work in building Distributions. This becomes more necessary as the number of Distributions for a Distributor grows.

Be sure to coordinate a Distribution's Build schedule with its Distributor's Refresh schedule and any related Channels' Send schedules. A Distribution build should begin after the Refresh schedule ends and finish before the Send schedules begin.

IMPORTANT: Do not refresh the Distributor more often than every five minutes. The following can need up to five minutes to complete their processes: Distribution building, eDirectory replication, and tree walking (when no Search policy is defined).

If you are using the Randomly Dispatch option, you should consider the End time for its Distributor's Refresh schedule when setting the Build schedule's Start time; and, you should consider the End time for the Build schedule when setting the Start times for the Send schedules.

Using the Randomly Dispatch Option in a Channel's Send Schedule

You can use the Randomly Dispatch option for a Channel's Send schedule to begin sending its Distributions to Subscribers randomly within a scheduling window. Each Distributor that has Distributions in the Channel calculates a random time between the specified Start and End times to begin sending its Distributions. This helps to balance the distribution workload for the network over a period of time.

For example, Distributor A and Distributor B have Distributions in a Channel. Each Distributor would calculate its own random time to begin sending its Distributions.

Another use of the Randomly Dispatch option for the Send schedule is if you have many Channels and you want all Distributions for all Channels to occur between 10 p.m. and 4 a.m. Using the Randomly Dispatch option in each Channel would allow you to disperse Distribution sending times for all Channels over that six-hour period of time.

If you are using the Randomly Dispatch option, you should consider the End time of each associated Distribution's Build schedule when setting the Send schedule's Start time; and, you should consider the End time for the Send schedule when setting the Start times for all associated Subscribers' Extract schedules.

Using the Randomly Dispatch Option in a Subscriber's Extract Schedule

You can use the Randomly Dispatch option for a Subscriber's Extract schedule to balance the Subscriber's work load in extracting Distributions.

If you are using the Randomly Dispatch option, you should consider the End times for the Send schedules of the Channels where the Subscriber is subscribed when setting the Start time for the Extract schedule.

Repeating Actions

For schedule types that have the Repeat the Action Every field, how this option works depends on other factors, such as other schedules and how frequently the Distributor reads eDirectory.

For example:

- ◆ You select Daily as the Send schedule for a Channel
- ◆ You set 1:00 a.m. to midnight (23 hours) as the sending window
- ◆ You set the Repeat the Action Every field with 1 hour as the repeat value

The action (sending the Distribution) repeats as follows:

1. Starting at 1:00 a.m. and repeating every hour, the Distributor queues the Distribution to be sent.
2. If a Distribution is in the process of being sent, it continues to be sent.
3. Once a Distribution is off the queue after being sent, the Distributor queues the next newer version for sending.

If a previously queued version of this Distribution has not been sent yet (still in the queue), the next newest version is placed in the queue. In other words, only one version of the Distribution (the last built) is queued while another version of the Distribution is being sent.

The Distributor always sends the latest Distribution, even if the Subscriber already has it.

8.3 Scheduling and Server Policies

All policies use the default schedule (Package Schedule) except where you change the schedule in a particular policy. You can also edit the default package schedule.

Review the following sections:

- ◆ [Section 8.3.1, “Policy Schedules Versus Distribution Schedules,” on page 338](#)
- ◆ [Section 8.3.2, “Scheduling a Server Policy,” on page 339](#)
- ◆ [Section 8.3.3, “Editing the Default Package Schedule,” on page 339](#)

8.3.1 Policy Schedules Versus Distribution Schedules

Policies that must be scheduled have two scheduling methods:

- ◆ Individual policy schedules, which are configured in a policy’s properties
- ◆ The Default Package Schedule that applies to all policies that are enabled in the package that do not have individual schedules set

With reference to policies that must be distributed, policy schedules and Distribution schedules are used for different purposes:

- ◆ **Enforcement:** A policy’s schedule determines when the policy can be enforced.
- ◆ **Distribution:** A Policy Package Distribution’s schedule determines when the policy’s Distribution is built so that the policy package can be sent, received, and extracted. After extraction, the policy’s schedule then determines when the policy can be enforced.

In other words, a Distribution's schedule does not directly determine when a policy is enforced. However, a Distributions's schedule, combined with the Channel's Send and Subscribers' Extract schedules, could delay the enforcement of a policy.

8.3.2 Scheduling a Server Policy

To schedule an individual policy:

- 1 In ConsoleOne, right-click a Policy Package object, click Properties, then select the Policies tab.
- 2 Select a policy, click Properties, then select the Policy Schedule tab.
- 3 Select a schedule in the Schedule Type field, then configure the schedule:

[Section B.1, "Daily," on page 400](#)

[Section B.2, "Event," on page 400](#)

[Section B.3, "Interval," on page 400](#)

[Section B.4, "Monthly," on page 401](#)

[Section B.5, "Never," on page 401](#)

[Section B.6, "Package Schedule," on page 401](#)

[Section B.7, "Relative," on page 402](#)

[Section B.8, "Run Immediately," on page 402](#)

[Section B.9, "Time," on page 402](#)

[Section B.10, "Weekly," on page 403](#)

[Section B.11, "Yearly," on page 403](#)

IMPORTANT: The Relative and Run Immediately schedules are not available for the Scheduled Down policy.

8.3.3 Editing the Default Package Schedule

To edit the default package schedule:

- 1 In ConsoleOne, right-click a Policy Package object, then click Properties.
- 2 Click Edit.
- 3 Select a schedule in the Schedule Type field, then configure the schedule:

[Section B.1, "Daily," on page 400](#)

[Section B.2, "Event," on page 400](#)

[Section B.3, "Interval," on page 400](#)

[Section B.4, "Monthly," on page 401](#)

[Section B.6, "Package Schedule," on page 401](#)

[Section B.7, "Relative," on page 402](#)

[Section B.8, "Run Immediately," on page 402](#)

[Section B.9, "Time," on page 402](#)

[Section B.10, "Weekly," on page 403](#)

[Section B.11, "Yearly," on page 403](#)

Review the following sections for information on variables in Novell® ZENworks® Server Management:

- ◆ Section 9.1, “Understanding Variables,” on page 341
- ◆ Section 9.2, “Types of Variables,” on page 344
- ◆ Section 9.3, “Defining a Variable,” on page 346
- ◆ Section 9.4, “Viewing All Variables in iManager,” on page 348
- ◆ Section 9.5, “Using a Variable to Change a Subscriber’s Console Prompt,” on page 348
- ◆ Section 9.6, “Using Variables to Control File Extraction,” on page 349

9.1 Understanding Variables

Review the following:

- ◆ Section 9.1.1, “Why Variables?,” on page 341
- ◆ Section 9.1.2, “Variable Usage,” on page 342
- ◆ Section 9.1.3, “Variable Usage Differences,” on page 343
- ◆ Section 9.1.4, “Precedence for Determining Which Variable to Use,” on page 344
- ◆ Section 9.1.5, “Distribution Variable Example,” on page 344

9.1.1 Why Variables?

You can use variables in Server Management to save time by more easily managing varying path information. For example, to globally control changes to the same location on all servers, you can use a variable for all volumes or drives in the script:

- ◆ Create the variable in each server’s Subscriber object where the value of the variable is the server’s volume or drive where the location exists.
- ◆ When the script runs, it passes the variable to each server, which in turn determines the variable’s value from the variable definition in its Subscriber object’s properties.
- ◆ The value identifies which volume or drive contains the desired location.

Using a variable for this information, you didn’t have to individually list each server’s name with its volume or drive in the script.

Variables are used to simplify referencing something that is specific to individual servers or software run on servers. For example:

| | |
|------------------------------------|--------------------------------|
| Destination Volumes or Drives | Script contents to be executed |
| DNS Names | Server Names |
| IP Addresses | Working Directories |
| Names of text files to be modified | |

Each of these can have different data per server. Variables allow you to account for those differences easily.

9.1.2 Variable Usage

- ♦ “Variable Syntax” on page 342
- ♦ “Nested Variables” on page 342
- ♦ “Literal % Symbols” on page 343

Variable Syntax

Variables can be thought of as having three parts: name, value, and usage. The syntax for each is:

Name syntax: *variable_name*

Example: DEST

Value syntax: *value_of_variable*

Example: sys:

Usage syntax: *%variable_name%*

Example: %DEST%

Thus, the variable DEST would equate to sys: on the particular server where the variable is defined.

When defining a variable, you do not provide the % character for the variable’s name or value. However, when using the variable, you use the % character before and after its name.

The software uses the % character to identify the beginning and ending of variable names. For example:

1. %DEST% tells the software that DEST is a variable name.
2. The software looks up DEST in a variable definition table on the receiving server to discover its value.
3. The value is then used to complete the path.

Nested Variables

You can nest variables to any level. For example, you can do the following to automate destinations:

1. Define DEST as the destination volume and directory:
 - ♦ **Variable name:** DEST
 - ♦ **Value:** %VOL_DRV%%DIR%
Here, you have nested two other variables inside of DEST to establish its value.
 - ♦ **Usage:** %DEST%
2. On a NetWare Subscriber, define the VOL_DRV variable:
 - ♦ **Variable name:** VOL_DRV
 - ♦ **Value:** *attribute*
For example, data:.

- ♦ **Usage:** %VOL_DRV%
3. On a Windows Subscriber, define the VOL_DRV variable:
 - ♦ **Variable name:** VOL_DRV
 - ♦ **Value:** *attribute*
For example, C:.
 - ♦ **Usage:** %VOL_DRV%
 4. Define the DIR variable:
 - ♦ **Variable name:** DIR
 - ♦ **Value:** *attribute*
A directory, such as \apps.
 - ♦ **Usage:** %DIR%

The result is that you can define the destination as simply DEST, which resolves to the directory and volume or drive specified at each target server. For example:

NetWare Subscriber: data:\apps

Windows Subscriber: c:\apps

Literal % Symbols

The % symbol is a valid character for file and directory names. Therefore, you need to identify literal usage of a % character. Otherwise, the software would think a nested variable name was being provided.

Literal % characters are identified by adding an extra % character immediately before a % character in the variable's value. This makes the software recognize the % character as a literal character and not a variable indicator. For example:

Variable name: DEST

Path for the variable: temp%abc%xyz

Variable value: temp%%abc%%xyz

The first % lets the software know that the next % character is literally part of the pathname, and not an indicator that a nested variable name is next. Without the double % characters, "abc" would be interpreted as a nested variable name.

9.1.3 Variable Usage Differences

General variable definitions, such as those in the Tiered Electronic Distribution policy, provide default variable values for Subscribers where they have none defined. Variables defined in a Subscriber object override such default variable values.

For Server Software Packages, variable names are resolved differently:

1. Is the variable defined in the Server Software Package component? If so, use that value.

IMPORTANT: A variable defined in a software package overrides any value defined in the Subscriber.

2. Is the variable one of the predefined variables? If so, use that value.
3. Is the variable a Java environment variable? If so, use that value.

9.1.4 Precedence for Determining Which Variable to Use

Variables are checked for in a specific order to determine which variable to use. The order is:

1. Server Software Packages ¹
2. Subscriber objects ¹
3. Tiered Electronic Distribution policy ¹
4. Default variables ²
5. Environment variables ²

¹ User-defined in Server Management

² Predefined

The variable is used from the first place where it is found.

9.1.5 Distribution Variable Example

Variables can also be used to specify where a Distribution is to be extracted, including the full path.

For example, you have a single Distribution with 20 Subscribers. You want to extract the Distribution to a specific volume on each of the Subscriber's servers. However, the volume name varies from server to server: 15 servers are using the data: volume and five are using voll:.

You can edit the Distribution Volume variable for some of these Subscribers by changing the Resolve To field on the Subscriber from data: to voll: for the five Subscribers using that volume.

When the Distribution is extracted, it goes to the correct volumes on each of the 20 servers.

9.2 Types of Variables

There are two types of variables:

- ♦ [Section 9.2.1, "Predefined Variables," on page 344](#)
- ♦ [Section 9.2.2, "User-Defined Variables," on page 346](#)

9.2.1 Predefined Variables

Predefined variables are created when ZENworks Server Management starts. They are used in Server Software Packages and Tiered Electronic Distribution, and are recognized by policy packages.

Predefined variables are not case sensitive, although they are displayed in all uppercase on the server console and in this documentation.

Syntax:

`%predefined_variable_name%`

where *predefined_variable_name* is the name defined by Server Management, and the % symbols tell the software that a variable name exists between them. For example:

`%WORKING_PATH%`

To make a predefined variable useful, its value must be set in the Server Software Package component, or in a Tiered Electronic Distribution object.

The Java environment can use predefined variables, such as `SERVER_DN` being used in a Java process call in a `.ncf` file.

An example of how a policy package can use a predefined variable is for the Broadcast Message text in the Server Down Process policy. The text can include a variable for the server name (`%SERVER_DN%`) so that the broadcast message displays the name of the server.

The Server Management predefined variables listed in [Table 9-1](#) are available:

Table 9-1 *Predefined Variables*

| Variable | Description and Value |
|--------------|--|
| BASE_PATH | Location of the Policy Manager: <code>sys:\zenworks\pds\smanager\</code> |
| CONF_PATH | Location of configuration files: <code>sys:\zenworks\pds\ted\</code> |
| IP_ADDRESS | IP address of a server, such as: <code>192.68.1.255</code> |
| LOAD_DIR | (NetWare® only) Directory where the server was loaded from: <code>c:\nwserver</code> |
| LOG_PATH | Location of log files: <code>sys:\zenworks\pds\smanager\</code> |
| PLUGINS_PATH | Where the Server Management plug-ins were installed: <code>sys:\zenworks\pds\smanager\plugins\</code> |
| POLICY_PATH | Where the policy files (.pol) are stored: <code>sys:\zenworks\pds\smanager\policy\</code> |
| PROP_PATH | Where Novell eDirectory™ object properties are stored: <code>sys:\zenworks\pds\smanager\prop\</code> |
| SERVER_DN | Distinguished server name in eDirectory, such as: <code>server01.servers.novell</code> |
| SERVER_NAME | Name given the server when NetWare was installed, such as: <code>server01</code> |

| Variable | Description and Value |
|--------------------|---|
| TED_PATH | Path to the ...ted directory: sys:\zenworks\pds\ted\ |
| TREE_NAME | Name of the eDirectory tree where Server Management servers reside. This is established during installation. |
| VOL | Default volume: sys: |
| WORKING_PATH | Working directory for the Server Policies and Server Software Packages components: sys:\zenworks\pds\smanager\working\ |
| ZWS_PATH | Where the ZENworks Web Server files are located: sys:\zenworks\zws\ |
| ZWS_PROP_FILE_PATH | Where the ZENworks Web Server property files are located: sys:\zenworks\zws\ |
| ZWS_SECURITY_PATH | Where the ZENworks Web Server security files are located: sys:\zenworks\zws\security\ |

9.2.2 User-Defined Variables

User-defined variables are created in the Server Software Package component, Subscriber objects, and the Tiered Electronic Distribution policy. Policy packages do not recognize user-defined variables.

User-defined variables are not case sensitive.

Syntax: `%variable_name%`

where *variable_name* is the name you give the variable when you define it. Spaces cannot be used in variable names. Use hyphens (-) or underscores (_) to separate words.

Variables defined in the Subscriber object are simple text substitutions. Text entered for the value of the variable is substituted for the variable name.

9.3 Defining a Variable

You can create variables in three locations:

- ◆ [Section 9.3.1, “Defining Default Variables for All Subscribers,” on page 347](#)
- ◆ [Section 9.3.2, “Defining Variables for a Specific Subscriber,” on page 347](#)
- ◆ [Section 9.3.3, “Defining Variables for a Server Software Package,” on page 348](#)

9.3.1 Defining Default Variables for All Subscribers

You can use the Tiered Electronic Distribution policy to define default variables for all Subscribers. Any variables you set in this policy as defaults for all Subscribers are overridden by any same-named variables defined on the Subscriber (see [Section 9.3.2, “Defining Variables for a Specific Subscriber,” on page 347](#)).

To define default variables:

- 1 In ConsoleOne[®], right-click a Service Location Package object, click *Properties*, select the check box for the Tiered Electronic Distribution policy to both select and enable it, click *Properties*, then select the *Variables* tab.
- 2 Click *Add*.
- 3 Provide the name of the variable.
The name can be user-defined, an environment variable (Java or native), or a predefined variable.
- 4 Provide the value for the variable.
The value is what the variable resolves to. It can also be another variable for nesting variables.
To ensure that extraction takes place, provide an absolute path to all Subscribers. For example, if the path is only the data: volume, make sure the colon (:) is included, because it is a necessary part of the full path.
- 5 Provide a description (optional), then click *OK*.
- 6 Repeat [Step 2](#) through [Step 5](#) to define another variable for the Subscribers.
- 7 Click *OK* when you have finished defining the default variables.
- 8 On the Service Location Package properties page, select the *Associations* tab.
- 9 If there are no associations listed that include all Subscriber objects, click *Add*.
- 10 Browse for an eDirectory container that includes all Subscriber objects, click *OK*.
This ensures that the policy is enforced for all Subscribers. For more information, see [“Tiered Electronic Distribution” on page 209](#).
- 11 Click *OK* to exit the policy package’s properties.

9.3.2 Defining Variables for a Specific Subscriber

- 1 In ConsoleOne, right-click a Subscriber object, then select the *Variables* tab.
- 2 Click *Add*.
- 3 Provide the name of the variable.
The name can be user-defined, an environment variable (Java or native), or a predefined variable.
Enter only the variable’s name. Do not include the % symbols that would accompany the variable when you use it.
- 4 Provide the value for the variable.
The value is what the variable resolves to. It can also be another variable for nesting variables.

To ensure that extraction takes place, provide an absolute path to the Subscriber. For example, if the path is only the data: volume, make sure the colon (:) is included, because it is a necessary part of the full path.

- 5 Provide a description (optional), then click *OK*.
- 6 Repeat **Step 2** through **Step 5** to define another variable for this Subscriber.
- 7 Click *OK* to exit the Subscriber's properties.

9.3.3 Defining Variables for a Server Software Package

- 1 In ConsoleOne, right-click a software package, then select the *Variables* tab.
- 2 Click *Add*.
New Variable #1 is defaulted in the *Variables* column.
Enter only the variable's name. Do not include the % symbols that would accompany the variable when you use it.
- 3 To provide a different name for the variable, use the Backspace key to delete the default name, type a new variable name, then press the Tab key.
The name can be user-defined, an environment variable (Java or native), or a predefined variable.
- 4 Provide the value for the variable.
The value is what the variable resolves to. It can also be another variable for nesting variables.
- 5 Repeat **Step 2** through **Step 4** to define another variable.
- 6 Click *OK* to exit the Server Software Package's properties.

9.4 Viewing All Variables in iManager

In Novell iManager, you can view all of the variables that are being used by a Server Management object:

- 1 In iManager, under the *ZENworks Server Management* role, click *Remote Web Console*.
- 2 Select a Distributor or Subscriber object, then click *OK*.
- 3 In the *Display* list box, select *Policy/Package Agent*.
This automatically displays the *Configurations* tab.
- 4 Scroll down to view the variables listed to the right of *Variables*.
All variables that can be used by the object are listed.

9.5 Using a Variable to Change a Subscriber's Console Prompt

The Subscriber can use the value of the PROMPT variable as its server console prompt.

- 1 In ConsoleOne, right-click a Subscriber object, then click *Properties*.
- 2 Select the *Variables* tab, then click *Add*.
- 3 In the Variables dialog box, provide information for the following fields:
Variable: Enter PROMPT as the variable name.

Value: Type the prompt text to be displayed. For example, %SERVER NAME% Subscriber could display as:

```
Provo_01 Subscriber >
```

Description: Provide a meaningful note (optional).

- 4 Click *OK* twice.

9.6 Using Variables to Control File Extraction

You can use variables to control the location that files are extracted on the Subscriber. Any destination can be used as a variable defined in a Subscriber object by encapsulating it with the percent (%) symbol.

IMPORTANT: Any variable value specified in the Tiered Electronic Distribution policy is a default value and is overridden by variable values set in a Subscriber object.

For the location where files are extracted, the destination root is identified in the File Grouping dialog box as a directory named `destrout`. This is the top-level directory used by a Subscriber to determine where to extract the file. The dialog box lets you build groups of directories under the `\destrout` directory.

You can specify the destination root as a known location (for example, %APP_DIR%). You can then go to the Variables tab on the Subscriber object and specify a value for this variable.

For example, variable APP_DIR would have the value:

```
sys:\apps
```

To use a variable to set the location that files are extracted to:

- 1 In ConsoleOne, right-click the Subscriber object, then click *Properties*.
- 2 Select the *Variables* tab, then click *Add*.
- 3 Provide the name of the variable.

The name can be user-defined, an environment variable (Java or native), or a predefined variable.

- 4 Provide the value for the variable.

The value is what the variable resolves to. It can also be another variable for nesting variables.

To ensure that extraction takes place, provide an absolute path to the Subscriber. For example, if the path is only the data: volume, make sure the colon (:) is included, because it is a necessary part of the full path.

- 5 Provide a description (optional), then click *OK* twice to exit the properties.

- 6 Create a new Distribution object.

For information, see [Section 3.4.4, “Creating a Distribution,” on page 123](#).

- 7 In the Distribution object’s properties, select the *Type* tab, select *File in the Select Type* drop-down box, then click *New Target*.

- 8 Replace the default %DEST_VOLUME% with the variable name, then click *OK* as necessary to exit the properties.

A directory named `\dest_volume` is created by default in the *Destination* column. You should select this directory to change the destination root. To select it, select the actual

directory name (destroot). You can then specify a known location or use a variable with surrounding percent symbols.

The following sections provide information for understanding and using the Novell® ZENworks® Server Management database in Policy and Distribution Services:

- ◆ [Section 10.1, “Understanding the ZENworks Database,” on page 351](#)
- ◆ [Section 10.2, “Determining How Many Databases You Need,” on page 353](#)
- ◆ [Section 10.3, “Installing and Connecting to the Server Management Database,” on page 356](#)
- ◆ [Section 10.4, “Creating a ZENworks Database Object,” on page 359](#)
- ◆ [Section 10.5, “Purging the Database,” on page 359](#)

10.1 Understanding the ZENworks Database

The following sections provide an understanding of the ZENworks database:

- ◆ [Section 10.1.1, “The Database Engine,” on page 351](#)
- ◆ [Section 10.1.2, “The Database File,” on page 351](#)
- ◆ [Section 10.1.3, “The Database Object,” on page 351](#)
- ◆ [Section 10.1.4, “Running the Database,” on page 352](#)
- ◆ [Section 10.1.5, “Database Caching,” on page 352](#)
- ◆ [Section 10.1.6, “Database Information,” on page 352](#)
- ◆ [Section 10.1.7, “Coexisting Databases,” on page 353](#)

10.1.1 The Database Engine

ZENworks Server Management is shipped with the Sybase database engine. This can only be installed once on a server. However, you can install Sybase to multiple servers.

Oracle* and SQL are not supported.

10.1.2 The Database File

Policy and Distribution Services uses a Sybase database file named `zfslog.db`. Server Management can function normally without the database, because it uses `zfslog.db` only to log information for Policy and Distribution Services reporting.

`zfslog.db` is normally located in the `\zenworks\pds\db` directory on a server. Its location is determined when using the installation program. It can reside on both NetWare® and Windows servers.

10.1.3 The Database Object

A Novell eDirectory™ database object is created during installation, named Server Management Database_ *server_name*. In its properties, the location of the database file (`zfslog.db`) is listed, if established during installation; otherwise, you can configure the ZENworks Database policy

(Service Location Package) to specify the database object. The location and policy are necessary for the database file to be found for logging information.

10.1.4 Running the Database

On NetWare servers, the database is run by using the `mgmt dbs.ncf` file (located in the `sys:\system` directory), which is executed from `autoexec.ncf`.

On Windows servers, the database is run by using the Novell Database - Sybase service.

10.1.5 Database Caching

Database files can become very large, which is why a 32 MB cache is recommended on the server where you are running the database. Caching improves server performance because of how frequently information is logged to `zfslog.db`.

10.1.6 Database Information

`zfslog.db` is used by Policy and Distribution Services to log successes and failures for the Server Policies or Tiered Electronic Distribution components. You can purge policy information automatically according to a policy setting. You can purge Tiered Electronic Distribution information manually from the database object. For information on purging, see [Section 10.5, “Purging the Database,” on page 359](#).

`zfslog.db` does not contain any configuration information.

The information listed in [Table 10-1](#) is written to `zfslog.db` by the agents:

Table 10-1 Agents that Write to the Database

| Agent | Information |
|----------------|---|
| Policy/Package | Failed and successful policies Discovered and unenforceable policies Down Server policy status Server Software Packages and components |
| Distributor | Distribution status: <ul style="list-style-type: none">◆ When built, sent, and extracted◆ Successes (plus reasons) of builds and extractions◆ Failures (plus reasons) of a build, send, receive, and extraction Subscriber status Revision histories |

For information on obtaining reports on the database information, see [Chapter 11, “Reporting,” on page 363](#).

The following provides information on gathering data for the database:

- ◆ A Distributor keeps track of each Subscriber in its routing hierarchy, so it knows which parent Subscribers have received a Distribution.
- ◆ The Distributor knows which Subscribers are at the end of a particular route, so it can know if Subscribers have not received a Distribution because a Subscriber higher up in the hierarchy failed to receive the Distribution.
- ◆ Subscribers send messages directly to the Distributor indicating that they have received a Distribution. The Distributor does not return a confirmation that it received the Subscriber's message.
- ◆ If a Distributor is not running when a "Successfully Received" message is sent from a Subscriber, this information is not written to the database. Because a message receipt confirmation is not received by the Subscriber, it does not re-send the message.

10.1.7 Coexisting Databases

You can have multiple Server Management databases in the tree. The number you have depends on whether you want consolidated reporting and can live with the additional network traffic in a WAN environment.

If you do not require consolidated reports, you can install one database file and object on different servers for each of your WAN segments. This eliminates writing to the database file over a WAN link by the Distributor.

For the server selected for a database file, you should not install a ZENworks Desktop Management database when a ZENworks Server Management database exists for Policy and Distribution Services. The Desktop Management database file replaces the ZENworks Server Management database file, causing all ZENworks Server Management database information to be lost. However, you can install a ZENworks Server Management database where a Desktop Management database exists and not lose any Desktop Management database information.

However, the databases for Management and Monitoring Services, Server Inventory, and Policy and Distribution Services can coexist on a server, because their database files use different filenames. You only need to name the database objects differently from each other, because they all have the same default object name of ZENworks Database.

10.2 Determining How Many Databases You Need

You can install the database to both NetWare and Windows servers.

The installation program checks the version of the Sybase engine before updating it. If it doesn't exist, or is an older version, Sybase software is installed.

IMPORTANT: Make sure you select a server for the database where you are installing the Subscriber/Policies option. The Purge Database option in the ZENworks Server Management policy (Distributed Server Package) works only if the Policy/Package Agent software and the `zfslog.db` file are located on the same server.

The installation program automatically creates a database object for each instance of the database that is installed, and you can select a database for the object during installation. You can install only

one instance of the database per run of the installation program. The database object is installed to the same eDirectory container as the Server object for the server where the database file, `zfslog.db`, is also installed.

Review the following to understand whether to have multiple database files:

- ♦ [Section 10.2.1, “Database Logging and Tiered Electronic Distribution Reporting,” on page 354](#)
- ♦ [Section 10.2.2, “Multiple Databases,” on page 355](#)

10.2.1 Database Logging and Tiered Electronic Distribution Reporting

Policy and Distribution Services can function normally without using a Server Management database, because it uses the `zfslog.db` file to only log information for reports. `zfslog.db` for Policy and Distribution Services does not contain any configuration information. To obtain Distribution status information, use the Tiered Distribution View and the Remote Web Console options under the *ZENworks Server Management* role in Novell iManager. Policy information (written to the database file) can be obtained through the canned reports available from the Tools menu in Novell ConsoleOne®.

The Distributor Agent writes its distribution status information (built, sent, received, extracted) and Server Software Package installation information to the database file (`zfslog.db`). The Policy/Package Agent writes policy enforcement successes and failures to the database file. This database information is used for the Policy and Distribution Services reports.

For the agents to know which database file they should write to, policies must be created for them. If not established during installation, the ZENworks Database policy can be used to associate (using the Service Location Package) so that the Distributor Agent can know where to write. The ZENworks Database policy is distributed (using the Distributed Server Package) to the Subscriber for the Policy/Package Agent to know where to write.

Policy and Distribution Services provides six predefined reports for the Server Policies component and four for the Tiered Electronic Distribution component. The report information is obtained from information logged to its database file. The reports listed in [Table 10-2](#) are available:

Table 10-2 *Policy and Distribution Services Reports*

| Server Policies Reports | Tiered Electronic Distribution Reports |
|-------------------------|--|
| Discovered Policies | Distribution Detail |
| Down Server Policy | Revision History |
| Packages | Revision History Failure |
| Failed Policies | Subscriber Detail |
| Successful Policies | |
| Unenforceable Policies | |

A selected report displays all of the applicable Server Policies or Tiered Electronic Distribution information currently logged in the database. The criteria you can specify for a report include date ranges, specific Distributions, Distribution versions, and so on.

You might want multiple databases for specialized reporting. For more information, see [“Advantages” on page 355](#).

For information on reporting, see [Chapter 11, “Reporting,” on page 363](#).

10.2.2 Multiple Databases

Policy and Distribution Services supports multiple instances of the Server Management database per tree. However, we recommend that you install only one instance of the database per tree. Review the following:

- ◆ [“Advantages” on page 355](#)
- ◆ [“Distributor Object Contexts and Multiple Databases” on page 355](#)
- ◆ [“Determining Whether You Need Multiple Databases” on page 356](#)

Advantages

The advantage in having only one database is that the Distribution information provided by all of the Distributor Agents and Policy/Package Agents can be displayed in a single report.

For example, with a single database, your software package information can be contained in one report:

- ◆ The Distributor Agent’s information on building and sending the Software Package Distribution
- ◆ The Policy/Package Agent’s information on extracting and installing the software package

The advantages in having multiple databases are:

- ◆ Minimizing traffic over slow WAN links

For example, having a separate database for Policy/Package Agent logging on its server’s side of a WAN link.

- ◆ Providing individual databases for specialized reporting

For example, if you have one database for the Distributor Agent (distributions) and one for the Policy/Package Agent (policies), the build and send information for the Software Package and Policy Package types of Distributions is written to the distributions instance of the database, and the software package installation and policy enforcement information is written to the policies instance of the database.

Distributor Object Contexts and Multiple Databases

One `zfslog.db` file can receive log entries from multiple Distributors, and a Distributor can only log to one `zfslog.db` file. The following explains why:

- ◆ For a Distributor Agent to locate a database file, it must have a ZENworks Database policy (Service Location Package) associated with a context above the Distributor’s object that points to the Database object, which contains the file’s location in its properties. (Distributors receive their policies through association.)
- ◆ If you have separate databases installed on two or more of your Distributor servers, each database requires its own ZENworks Database policy for locating it (the policy points to the database’s object, which contains its file’s location).

- ◆ Only one Service Location Package (which contains the ZENworks Database policy) can be associated with a given context, such as the container holding your Distributor objects.
- ◆ Because only one Service Location Package can be associated with a given context, you must install your Distributor objects to different contexts to have multiple Distributors writing to their individual database files. Each Distributor would need its own database location policy that is associated with its own parent container.

For ease of management, you can keep your Distributor objects near each other by creating individual containers for each of them under the container where you usually place all of them. Then you can associate the different Service Location Packages with their appropriate Distributor's unique parent containers.

- ◆ To have all of your Distributors write to the same database file, place each of their Distributor objects somewhere under the container where you associate the Service Location Package. They would all use the same database location policy.

Determining Whether You Need Multiple Databases

Consider the following to determine how many databases to have in the tree:

- ◆ **WAN traffic:** Tiered Electronic Distribution does not perform a large number of database updates, so the actual impact on system resources should be minimal. The greatest impact could be the time it takes to perform the transaction. However, if you have slow WAN connections, you might not want database logging to occur over the WAN.
- ◆ **Multiple Distributors:** If you have multiple Distributors in the tree, you can have one database for each, or have them share one or more databases. The type of Distributor reporting you want should determine whether to have a separate database for each. For example, are your Distributors specialized in the types of Distributions they'll send?
- ◆ **Consolidated reporting:** To have only one report for all of your Tiered Electronic Distribution information, install only one database object and file and have all Distributors log to that one file, regardless of WAN traffic considerations. Use the ZENworks Database policy (Service Location Package) to direct all Distributors to that database file.
- ◆ **Specialized reporting:** You might want reports that are specific to a region or group of servers. You can install a database object and file for each such region and have the Distributors in those regions or server groups log to that database. Use a separate ZENworks Database policy (Service Location Package) to direct each Distributor to its desired database file.

10.3 Installing and Connecting to the Server Management Database

You should install the Server Management database on a server where policies are enforced. This is required so that you can use the ZENworks Database policy to locate the database file, `zfslog.db`.

The Server Management Database object is automatically created in the tree when you run the installation program and select a server for the database.

The installation program can install only one database at a time. To install additional databases to the tree, you need to perform the steps in the following sections for each database to be installed.

Perform the steps in the following sections to install and set up the database:

- ♦ [Section 10.3.1, “Installing the Database,” on page 357](#)
- ♦ [Section 10.3.2, “Connecting to the Database,” on page 358](#)

10.3.1 Installing the Database

To install a Policy and Distribution Services database:

- 1 On a workstation, insert the *ZENworks 7 Server Management with Support Pack 1 Program* CD.

The startup screen is displayed. If the startup screen is not automatically displayed after inserting the CD, you can start it by running `winsetup.exe` at the root of the CD.

IMPORTANT: Installation from a remote CD is not supported unless there is a drive mapped on the workstation to that CD. For example, if you place the CD in a Windows server CD drive, then run the installation from a workstation, you must have a drive mapped to the CD drive of that Windows server.

- 2 Select the *Server Management* option.
- 3 Click *Policy-Enabled Server Management* to start the installation program.
- 4 If you agree with the Software License Agreement, click *Accept*, then click *Next* to display the Installation Type page; otherwise, select *Decline* and click *Cancel* to exit.
- 5 On the Installation Type page, click *Next* to perform a new installation and display the Installation Options page.
- 6 On the Installation Options page, click *Next* to accept the defaults and display the eDirectory Tree for Creating Objects page.
- 7 Browse and select the tree to install to (you can only select one tree), then click *Next*.

The tree name is not case sensitive.

- 8 On the Server Selection page, click *Add*, then browse for the server where you want to install the database.

You can select only one server per run of the installation program.

You might want a database for each Distributor to write its own information to. However, Distributors can share a database. Because the Distributor writes information to the database for all Tiered Electronic Distribution objects, you should install the database on the same server as the Distributor to minimize network traffic.

IMPORTANT: Make sure you select a server for the database where you are installing policies. The Purge Database option works only if the `zfs.ncf` and `zfslog.db` files are on the same server.

- 9 Under *Additional Options*, select the *Server Management Database* check box to enable it, then click *Next* to display the File Installation Paths and Options page.

The installation program checks all mounted volumes on the server to see if `zfslog.db` exists. If not, both the file and the database object are installed. If the file exists, the database object is still installed.

- 10 Click *Next* to accept the defaults and display the Database Settings page.

- 11 To change the default path to the database file, edit the *Database Path* field.

IMPORTANT: Because the database file can become very large, we recommend that you change the default NetWare volume from sys: to another volume on that server.

- 12 Accept the other defaults on the Database Settings page by clicking *Next* to display the Policy and Distribution Services Database Logging page.
- 13 To determine logging for the Server Management database that you configured in a previous installation page, select one of the following:

Log to an existing Server Management database: Select an existing database file for logging by browsing for and selecting the database object to associate it with.

Log to a Server Management database that will be installed: The database object name that you configured in a previous installation page is the default. However, you can browse for and select an existing database object.

Do not log to a Server Management database: You can elect to not log to a database at this time, even though you have configured a database in the previous installation page.

- 14 On the Summary page, review your selections, then click *Finish*.

The installation program now copies files and installs the database objects.

WARNING: If you click Cancel, none of the work you did in the installation program is saved.

After the installation has finished, you can check the installation log file (see [Step 10](#)) to see if any components failed to install.

The ZENworks Database policy is automatically created and configured during installation of this new database.

- 15 Continue with [Section 10.3.2, “Connecting to the Database,”](#) on page 358.

10.3.2 Connecting to the Database

To make sure that the database can be written to by the Policy/Package Agent:

- 1 On a server, load the agent by doing the following:

| Server Platform | Agent Startup Method |
|------------------|---|
| Windows | 1. Open the Control Panel. 2. Click Admin Tools, then click Services. 3. Click Novell ZENworks Service Manager, then click Start. |
| NetWare | <code>sys:\zenworks\pds\smanager\zfs.ncf</code> |
| Solaris or Linux | <code>/etc/init.d/Novell-ZFS Start</code> |

Note whether a message is displayed indicating that the agent has connected to the database.

- 2 To determine whether the agent is writing to the database, do the following:

- 2a At a NetWare server’s console prompt, view the monitor while the agent is loading.

A message should display that states whether the agent connected with the database.

- 2b** If the message indicates that the agent did not connect to the database, you should check the following:
- ♦ Is the database is running on the server?
 - ♦ Is there a database object that has its Policy/Distribution Management tab set up with the server where the database file is installed?
 - ♦ Is there an effective ZENworks Database policy pointing to the database object?

10.4 Creating a ZENworks Database Object

The ZENworks Database object might not exist if you had inadvertently deleted the object.

If the database object does not exist in the tree because you didn't originally install it, you must use the GUI installation program to create it. For more information, see “[Starting the Installation Program](#)” in the *Novell ZENworks 7 Server Management Installation Guide*. Follow only the steps and select only the options that are necessary to create the database.

To re-create a database object that was inadvertently deleted:

- 1** In ConsoleOne, right-click a location in the tree for the database object, click *New > Object*, then click *ZENworks Database*.
- 2** Provide a database name.
- 3** Select the *Define Additional Properties* check box, then click *OK*.
- 4** On the *ZENworks Database* tab, select either the *Server DN* or *Server IP Address* option.
One of these location IDs could already be the default. If not, provide the information for the server where *zfslog.db* resides.
- 5** Select the *eDirectory Rights* tab, click *Trustees of This Object*, click *Add Trustee*, then select *[Public]*.
The database object must be assigned a trustee of Public, or the Policy/Package Agent displays messages that it cannot connect with the database or read the ZENworks Server Management policy.
- 6** Click *OK*.
If you click *Cancel*, none of the information you added or changed on any of the tabs is saved. However, the database object remains on the tree.
- 7** Set up the ZENworks Database policy.
For steps to specify the location of a database, see “[ZENworks Database](#)” on page 230.
- 8** Associate the Service Location Package with a container above where the Distributor object resides.

10.5 Purging the Database

Because Policy and Distribution Services logs all successes and failures for the Server Policies or Tiered Electronic Distribution components, *zfslog.db* can quickly grow in size. Therefore, you should periodically purge this database file.

The following database information types are purged using different methods:

- ♦ [Section 10.5.1, “Tiered Electronic Distribution Information,”](#) on page 360

- ◆ [Section 10.5.2, “Server Policies Information,” on page 360](#)

10.5.1 Tiered Electronic Distribution Information

To manually purge a selected database of all Tiered Electronic Distribution information older than a specific date and time:

- 1 In ConsoleOne, right-click the database object, then click *Purge*.
- 2 In the Purge Database dialog box, select a date and time, then click *OK*.

Records older than the date entered are purged from the database.

When the purge has been completed, a dialog box is displayed indicating that the purge was successful.

10.5.2 Server Policies Information

Purging of policy information is done automatically according to how you configure the ZENworks Server Management policy and which of the following events occurs:

- ◆ A server is restarted where the Policy/Package Agent is running that writes policy information to the database.
- ◆ Server Management is restarted on a server where the Policy/Package Agent is running that writes policy information to the database.
- ◆ On a server where The Policy/Package Agent is running that writes policy information to the database, the Policy/Package Agent is manually refreshed by typing the REFRESH command on the ZENworks Server Management console prompt.

The REFRESH command or Refresh option only causes database purging if given on a server where the database file resides.

In each of these events, the database file that is purged is the one written to by the associated Policy/Package Agent.

To set up policy information purging:

- 1 In ConsoleOne, do one of the following:
 - ◆ If you want to use a different policy package schedule for purging information than an existing Distributed Server Package (see [Step 5](#)) is using, in ConsoleOne click *File > New > Policy Package*, select *Distributed Server Package*, and provide a name that identifies its purpose. For example, *Purge_Policy_server_name*, where *server_name* is the server where the `zfslog.db` file resides.
 - ◆ To use the same policy package that has other policies enabled, in ConsoleOne right-click the existing Distributed Server Package, then click *Properties*.
In this case, we recommend that in [Step 5](#) you select *Run Immediately* for the package schedule. That way, any change you make to the number of days in [Step 3](#) is immediately available the next time a purge is triggered.
- 2 In the Distributed Server Package, select the *ZENworks Server Management policy* check box, then click *Properties*.
- 3 Select the *ZENworks Server Management Configuration* tab and select a number of days.

The default is 100 days. Records older than the number of days that you determine are purged.

Select a number that maintains the desired database file's size. The amount of policy-related information accrued in the database is determined by how often you have policies being run by servers writing to this database. Depending on how frequently you purge the database, you may need to experiment over time to determine the optimum number of days.

4 Click *OK* to close the policy's properties.

5 To set the package schedule, do one of the following:

- ♦ To accept the default package schedule, which is Run Event: System Startup, click *OK* to close the package's properties.
- ♦ To change the default schedule, click Edit, select a schedule, then click *OK* twice to close the package's properties.

For more information on the schedules, see [Section 4.7, "Scheduling Policies," on page 234](#).

The package schedule determines when any configuration changes that you make are available. For example, if you previously selected Event: System Startup for the package schedule and then later changed the 100 days to 60, that change is not recognized if the Policy/Package Agent is refreshed to trigger purging. It is only recognized after system startup occurs.

6 Create a Policy Package Distribution for this policy.

For more information, see ["Creating and Configuring the Distribution" on page 57](#).

7 Send the Distribution to the Subscriber server where the `zfslog.db` file resides.

For more information, see [Section 1.2.6, "Sending the Distributions," on page 60](#).

If this policy package is dedicated the ZENworks Server Management policy for purging, you need to send this Distribution only to each server where a database file resides, because you need just one instance of this policy per database file.

Novell® ZENworks® Server Management provides predefined reports for the Policy and Distribution Services components:

- ◆ [Section 11.1, “Understanding Policy and Distribution Services Reporting,” on page 363](#)
- ◆ [Section 11.2, “Report Descriptions,” on page 365](#)
- ◆ [Section 11.3, “Generating Reports,” on page 369](#)
- ◆ [Section 11.4, “Creating Customized Reports,” on page 370](#)

11.1 Understanding Policy and Distribution Services Reporting

Review the following:

- ◆ [Section 11.1.1, “Reporting Categories,” on page 363](#)
- ◆ [Section 11.1.2, “Reporting Scope,” on page 364](#)
- ◆ [Section 11.1.3, “Accessing Reports,” on page 364](#)
- ◆ [Section 11.1.4, “Creating and Storing Report Information,” on page 364](#)

11.1.1 Reporting Categories

Server Policies has six predefined reports, and Tiered Electronic Distribution has four. For details on each predefined report, see [Section 11.2, “Report Descriptions,” on page 365](#).

The following sections describe the purposes of the Policy and Distribution Services reports:

- ◆ [“Purposes for the Server Policies Reports” on page 363](#)
- ◆ [“Purposes for the Tiered Electronic Distribution Reports” on page 363](#)

Purposes for the Server Policies Reports

Server Policies reports show which servers have processed which policies, when they were processed, and if their enforcement was successful.

Purposes for the Tiered Electronic Distribution Reports

The Distribution-level reports show the view from the Distributor side and are very useful for checking which Subscribers succeeded or failed to receive and extract a particular Distribution. The Subscriber reports are used to determine which Distributions a single Subscriber has received.

Reporting gives very detailed information regarding which nodes succeeded. All known error conditions are caught and error conditions are reported to the database. However, when a process status is in progress, errors can occur or failures can occur on the node that are not caught (for example, the machine went down or the process was killed).

Subscribers that did not attempt to receive the Distribution (because they were not set up correctly or were not running) do not have information displayed on the report. You can compare the number expected against the actual numbers and look for missing Subscribers on the report. After Subscribers are set up and have been functioning, this should not be a common problem.

11.1.2 Reporting Scope

A selected report displays all of the applicable Server Policies or Tiered Electronic Distribution information currently logged in the database. There are options for defining the selection criteria for the data that will appear on some reports, such as date ranges, or for selecting Policy Package objects or Tiered Electronic Distribution objects.

11.1.3 Accessing Reports

There are two access points for Policy and Distribution Services reports:

Via the Object

- 1 In ConsoleOne, right-click a ZENworks Database object.
- 2 Click *Reporting*.

The report dialog box for the Policy and Distribution Services canned reports is displayed.

Via the Menus

- 1 In ConsoleOne, click *Tools > ZENworks Reports*.
- 2 Click *Reporting*.

The report dialog box for the Policy and Distribution Services canned reports is displayed.

11.1.4 Creating and Storing Report Information

A Policy and Distribution Services database file (`zfslog.db`) is used to store the report information. After you have installed and run the database and data has been placed in `zfslog.db`, Policy and Distribution Services reporting is enabled.

The Policy/Package Agent running on each Subscriber server writes Server Policies information to the database. The Distributor Agent writes the Tiered Electronic Distribution information and the Server Software Package information to the database.

Each Distributor may normally have its own ZENworks Database object and database file (`zfslog.db`), so report information could be given only for the particular Distributor associated with the ZENworks Database object selected.

Information is logged to the `Zfslog.db` file when any of the following actions have occurred:

- ♦ The ZENworks Database policy (Service Location Package) has been configured and enabled (see “ZENworks Database” on page 214)
- ♦ The ZENworks Database policy (Distributed Server Package) has been configured and enabled (see “ZENworks Database” on page 214)
- ♦ The Policy/Package Agent has either been refreshed from the server console or Server Management has been restarted

The ZENworks Database policy (contained in the Distributed Server Package) must already have been received and extracted on the Subscriber server before the Policy/Package Agent can log to the database file.

- ◆ The Distributor Agent has been restarted (not refreshed) after the ZENworks Database policy has been enabled (which includes associating it with the Distributor object's container).

11.2 Report Descriptions

The following sections describe the Policy and Distribution Services reports:

- ◆ [Section 11.2.1, “Tiered Electronic Distribution Reports,” on page 365](#)
- ◆ [Section 11.2.2, “Server Policy Reports,” on page 367](#)

11.2.1 Tiered Electronic Distribution Reports

There are four predefined Tiered Electronic Distribution reports:

- ◆ [“Distribution Detail Report” on page 365](#)
- ◆ [“Revision History Report” on page 365](#)
- ◆ [“Revision History Failure Report” on page 366](#)
- ◆ [“Subscriber Detail Report” on page 366](#)

Distribution Detail Report

Displays a detailed, time-line history of Distributions for the selected Subscribers, including:

- ◆ Distributions Sent
- ◆ Distributions Received
- ◆ Distributions Extracted (including start time, end time, and completion code)

Sorting is by time; grouping is by Distribution name and version.

The report criteria include:

- ◆ **Subscriber:** If you selected a Subscriber object, it appears in the Subscriber field and the report only displays information for the receive and extract actions performed by this Subscriber. Information for parent Subscribers also displays a Received Stage heading.

If you selected the Database or Distribution object, the report includes all actions that have occurred with a Distribution. In other words, information for all Subscribers involved is displayed.

- ◆ **Latest version only:** Deselect to include versions that are within the specified date range.
- ◆ **Select the date range criteria for the report:** Specify the range.

Revision History Report

Displays a history of a Distribution package's versions, including:

- ◆ Distribution (DN of package)
- ◆ Distributor (DN of object)

- ◆ Version Number
- ◆ Creation Date/Time
- ◆ Distribution Size

Sorting is by version number.

The report criteria include:

- ◆ **Distribution:** If you selected a Distribution object, it appears in the Distribution field. If you selected the Database object, you need to browse for the Distribution object.

Revision History Failure Report

Displays the versions of the Distribution that failed during creation, including:

- ◆ Distribution (DN of package)
- ◆ Distributor (DN of object)
- ◆ Creation Date and Time
- ◆ Error Description

Sorting is by version.

The report criteria include:

- ◆ **Distribution:** If you selected a Distribution object, it appears in the Distribution field. If you selected the Database object, you need to browse for the Distribution object.

Subscriber Detail Report

Displays status information for the Subscribers that received the Distribution, including:

- ◆ Distribution and Version
- ◆ Subscriber (DN of object) and Subscriber's Address
- ◆ Channel Name
- ◆ Source (DN of Distributor)
- ◆ Stage
- ◆ Status
- ◆ Date and Time
- ◆ Error Description

Sorting is by Subscriber/Parent Subscriber, then Stage.

The report criteria include:

- ◆ **Distribution:** If you selected a Distribution object, it appears in the Distribution field. If you selected the Database object, you need to browse for the Distribution object.
- ◆ **Version number:** If Distribution versions exist, you can choose one from the drop-down menu. Select All to include all versions.
- ◆ **Distribution stage:** You can select All, Extract, or Receive.
- ◆ **Distribution status:** You can select All, Success, or Not Success.

11.2.2 Server Policy Reports

For all server policy reports, the default date ranges are for the current date (from midnight to midnight).

There are six predefined server policy reports:

- ◆ “Discovered Policies Report” on page 367
- ◆ “Server Down Process Report” on page 367
- ◆ “Failed Policies Report” on page 368
- ◆ “Packages Report” on page 368
- ◆ “Successful Policies Report” on page 368
- ◆ “Unenforceable Policies Report” on page 369

Discovered Policies Report

Displays the servers that have discovered policies within the specified packages, including:

- ◆ Package (DN)
- ◆ Server DN
- ◆ Server Name
- ◆ OS Name and OS Version
- ◆ Date/Time of Discovery

Sorting is by package, then by context/server name, maintaining the tree’s hierarchy. For example, `myserver.servers.novell` is sorted `novell`, `servers`, `myserver`.

The report criteria include:

- ◆ **Package:** Select a policy package from the drop-down list or select All.
- ◆ **Policy type:** You can select All, Server Down Process, Scheduled Down, SNMP Trap Targets, Community Strings, Set Parameters, Script, Text File, Scheduled Load/Unload, or Database Location.
- ◆ **Select the date range criteria for the report:** Specify the range.

Server Down Process Report

For a selected server or all servers in the tree, displays Server Down Process policy information, including:

- ◆ Down Action and Code for each policy

Sorting is by server name only.

The report criteria include:

- ◆ **Server:** Select a server from the drop-down list or select All.
- ◆ **Select the Date Range Criteria for the Report:** Specify the range.

Failed Policies Report

For all servers in the tree, displays all policies that have failed, including:

- ◆ Package (DN)
- ◆ Server DN
- ◆ Server Name
- ◆ OS Name
- ◆ Date/Time of Failure
- ◆ Reason for Failure (Description)

Sorting is by context/server name, maintaining the tree's hierarchy. For example, `myserver.servers.novell` is sorted `novell`, `servers`, `myserver`.

The report criteria include:

- ◆ **Package:** Select a policy package from the drop-down list or select All.
- ◆ **Failure type:** You can select All, Failed, Unenforceable, or Partial Enforcement.
- ◆ **Policy type:** You can select All, Server Down Process, Scheduled Down, SNMP Trap Targets, Community Strings, Set Parameters, Script, Text File, Scheduled Load/Unload, or Database Location.
- ◆ **Select the date range criteria for the report:** Specify the range.

Packages Report

Displays information on Server Software Packages and their components, including:

- ◆ Success status of each package
- ◆ Success status of each component

Sorting is by context/server name, maintaining the tree's hierarchy. For example, `myserver.servers.novell` is sorted `novell`, `servers`, `myserver`.

The report criteria include:

- ◆ **Package:** Select a software package from the drop-down list or select All.
- ◆ **Server:** Select a server from the drop-down list or select All.
- ◆ **Select the date range criteria for the report:** Specify the range.

Successful Policies Report

For all servers in the tree, displays all policies that have been successfully enforced, including:

- ◆ Package (DN)
- ◆ Server DN
- ◆ Server Name
- ◆ OS Name
- ◆ Date/Time of Run
- ◆ Action Code

Sorting is by context/server name, maintaining the tree's hierarchy. For example, `myserver.servers.novell` is sorted `novell`, `servers`, `myserver`.

The report criteria include:

- ◆ **Package:** You can specify a single policy package or select All.
- ◆ **Success type:** You can select All, Change, or No Change.
- ◆ **Policy type:** You can select All, Server Down Process, Scheduled Down, SNMP Trap Targets, Community Strings, Set Parameters, Script, Text File, Scheduled Load/Unload, or Database Location.
- ◆ **Select the date range criteria for the report, from/to:** Specify the range.

Unenforceable Policies Report

Displays all unenforceable policies because of the absence of an enforcer on a server for all servers in the tree, including:

- ◆ Package (DN)
- ◆ Server DN
- ◆ Server Name
- ◆ OS Name and OS Version

Sorting is by package, then by server name.

The report criteria include:

- ◆ **Package:** Select a policy package from the drop-down list or select All.
- ◆ **Select the date range criteria for the report:** Specify the range.

11.3 Generating Reports

- 1 In ConsoleOne, right-click the ZENworks Database object.

The ZENworks Database object must be one where its Policy/Distribution Management tab (not the Inventory Management tab) is properly configured.

- 2 Click *Reports*.

- 3 Select a report:

- ◆ Server Policy Reports
 - Discovered Policies
 - Failed Policies
 - Packages
 - Server Down Process Policy
 - Successful Policies
 - Unenforceable Policies
- ◆ Tiered Electronic Distribution Reports
 - Distribution Detail
 - Revision History

Revision History Failure
Subscriber Detail

4 Select the reporting criteria.

If you need more detail on reporting criteria or content, see [Section 11.2, “Report Descriptions,” on page 365](#).

5 Click *Run Selected Report*.

The View Report dialog box is used to display the generated report. The dialog box has the following features:

- ◆ To expand how much of the report you can view on screen, resize the dialog box.

The report is displayed in landscape orientation for printing purposes.

- ◆ Use the following navigation options to move through the report:

First Page

Previous Page

Next Page

Last Page

Go To Page

- ◆ To print the report, click File > Print.
Your default printer is selected.
- ◆ To export the report, click File > Export Report.

You can export to the following formats:

Text

HTML

PDF

SDF

11.4 Creating Customized Reports

Using the following database information, you can create custom reports for the Server Policies and Tiered Electronic Distribution components.

However, for Tiered Electronic Distribution objects such as a Subscriber or the External Subscriber, you should use ZENworks reporting options (see [Chapter 11, “Reporting,” on page 363](#)) or iManager ([Chapter 2, “Novell iManager,” on page 63](#)) for determining the status of Distributions or policies.

The database file (`zfslog.db`) contains the following information:

- ◆ [Section 11.4.1, “Default Sybase Database User ID and Password,” on page 370](#)
- ◆ [Section 11.4.2, “Server Policies Database Contents,” on page 371](#)
- ◆ [Section 11.4.3, “Tiered Electronic Distribution Database Contents,” on page 378](#)

11.4.1 Default Sybase Database User ID and Password

The Sybase database (`zfslog.db`) that ships with Server Management has the following default user ID and password:

User ID: dba

Password: sql

11.4.2 Server Policies Database Contents

Following are the database table definitions for server policies:

- ♦ “SERVERS” on page 371
- ♦ “SERVERIP” on page 371
- ♦ “PACKAGES” on page 372
- ♦ “POLICIES” on page 372
- ♦ “POLICYACTION” on page 373
- ♦ “PACKAGEACTION” on page 374
- ♦ “SOFTWARECOMPONENTACTION” on page 376
- ♦ “Foreign Keys” on page 377

SERVERS

Contains one record for each server running the Policy/Package Agent.

Table 11-1 Servers Field Names

| Field Name | Type | Use |
|------------|---------|--|
| SERVID | integer | not null Unique number that is automatically assigned. |
| SERVERNAME | vchar | not null The short name of the server as seen on the console prompt. |
| SERVERDN | vchar | DN of the Server object in eDirectory (dot separated). |
| REVERSEDN | vchar | not null SERVERDN in reverse order and backslash (\) delimited. |
| OSNAME | vchar | Name of the operating system, such as NetWare 5.1. |
| OSVERSION | char | Version of the operating system, such as 5.1, 6.0, and so on. |
| TREENAME | vchar | Name of the eDirectory tree containing the server. |

Primary key (SERVERID)

SERVERIP

Contains one record for each server running the Policy/Package Agent.

Table 11-2 *ServerIP Field Names*

| Field Name | Type | Use |
|-------------|---------|--|
| SERVERIPKEY | integer | not null Assigned automatically: Default Auto increment. |
| SERVERID | integer | not null Links to the SERVERS table. |
| IPADDRESS | varchar | not null Server's IP address. |

Primary key (SERVERID) REFERENCES SERVERS
Primary key (SERVERIPKEY)

PACKAGES

Contains one record for each version of a software package that the Policy/Package Agent has attempted to process.

Table 11-3 *Packages Field Names*

| Field Name | Type | Use |
|----------------|---------|---|
| PACKAGEGUID | char | not null Assigned automatically. |
| PACKAGENAME | char | Name of .cpk file or policy package. |
| PACKAGEDESC | char | Description contained in a Server Software Package component. |
| PACKAGEVERSION | char | Version of the software package. |
| BUILDDATE | integer | Date the software package was compiled. |

Primary key (PACKAGEGUID)

POLICIES

Contains one record for each policy or policy package combination.

Table 11-4 *Policies Field Names*

| Field Name | Type | Use |
|---------------|---------|---|
| POLICYID | integer | not null A globally unique ID. |
| POLICYDN | varchar | The DN of the eDirectory policy object. |
| POLICYPACKAGE | varchar | The DN of the policy package the policy belongs to. |

| Field Name | Type | Use |
|----------------|---------|--|
| POLICYCLASS | varchar | The class or type of policy. For definitions, see “Valid Entries for POLICYCLASS” on page 373. |
| POLICYTREENAME | varchar | The name of the tree the policy object is in. |

Primary key (POLICYID)

Valid Entries for POLICYCLASS

zenZFSServerDowningPolicy
zenZFSScheduleDownPolicy
zenZFSSetServerParamPolicy
zenZFSServerScriptPolicy
zenZFSTextFilePolicy
zenZFSScheduledRunPolicy
zenZFSZFSPolicy
zenZFSCommunityPolicy
zenZFSSNMPTrapTargetPolicy
zenZFSSMTPHostPolicy
zenZFSDatabaseLocationPolicy
zenZFSLicenseLocationPolicy
zenZFSTEDPolicy

POLICYACTION

Contains one record for each action performed.

Table 11-5 PolicyAction Field Names

| Field Name | Type | Use |
|-----------------|-----------|--|
| POLICYACTIONKEY | integer | not null Assigned automatically: Default Auto increment. |
| POLICYID | integer | not null Links to the POLICIES table. |
| SERVERID | integer | not null Links to the SERVERS table. |
| CREATIONDATE | timestamp | Time stamp of the action. |
| DESCRIPTION | varchar | Undefined string describing an error. |
| CODE | integer | Code representing the result of the action. For definitions, see “Valid Entries for CODE” on page 374. |
| ACTIONCODE | integer | The action being performed. For definitions, see “Valid Entries for ACTIONCODE” on page 374. |

Primary key (POLICYACTIONKEY)

Valid Entries for CODE

| | |
|------------------------|------|
| RC_POL_SUCCESS | = 0 |
| RC_POL_PARTIAL_SUCCESS | = 1 |
| RC_POL_FAILURE | = -1 |
| RC_POL_EMPTY | = -2 |

Exception: If the value in the ACTIONCODE field is AC_POL_DOWN_CONNECTIONS or AC_POL_DOWN_DISCONNECTIONS, then the value of CODE is either the current number of active connections, or the number of forced disconnects.

A number 1 in the CODE field can mean one of the following:

- ◆ There was a partial success
- ◆ There is one active connection
- ◆ There was one forced disconnect

This is because the meaning of the entry in the CODE field is determined by the content of the ACTION CODE field.

Valid Entries for ACTIONCODE

| | |
|----------------------------|-------|
| AC_POL_DISCOVERED | = 101 |
| AC_POL_SCHEDULED | = 102 |
| AC_POL_APPLIED | = 103 |
| AC_POL_APPLIED_CHANGE | = 104 |
| AC_POL_NO_ENFORCER | = 105 |
| AC_POL_DOWN_CONNECTIONS | = 106 |
| AC_POL_DOWN_DISCONNECTIONS | = 107 |
| AC_POL_DOWN_UNLOAD | = 108 |
| AC_POL_DOWN_EMAIL | = 109 |
| AC_POL_DOWN_NOTIFY | = 110 |
| AC_POL_DOWN_CANCELED | = 111 |
| AC_POL_DOWN_IGNORED | = 112 |
| AC_POL_DOWN_REQUESTED | = 113 |

PACKAGEACTION

Contains one record for each action taken on a Server Software Package.

Table 11-6 PackageAction Field Names

| Field Name | Type | Use |
|------------------------|-----------|--|
| PACKAGEACTIONID | integer | not null Assigned automatically: Default Auto increment. |
| PACKAGEGUID | char | not null Links to the PACKAGES table. |
| SERVERID | integer | not null Links to the SERVERS table. |
| CREATIONDATE | timestamp | Time stamp of the action. |
| DESCRIPTION | varchar | For definitions, see “Valid Entries for DESCRIPTION” on page 375. |
| CODE | integer | Code representing the results of the action. For definitions, see “Valid Entries for CODE” on page 375. |
| ACTIONCODE | integer | Code representing the action being performed. For definitions, see “Valid Entries for ACTIONCODE” on page 375. |
| STARTEDPACKAGEACTIONID | integer | 0 = started running the package, or when the new action is logged then the PACKAGEACTIONID of the new action replaces the 0. |

Primary key (PACKAGEACTIONID)

Valid Entries for DESCRIPTION

Started package
 Finished rollback
 Error description
 Or it is empty

Valid Entries for CODE

| | |
|---------|-----|
| Success | = 0 |
| Failure | = 1 |
| Partial | = 2 |

Valid Entries for ACTIONCODE

| | |
|-----------------------------|-----|
| AC_PACKAGE_INSTALL | = 0 |
| AC_PACKAGE_ROLLBACK | = 1 |
| AC_PACKAGE_INSTALL_STARTED | = 2 |
| AC_PACKAGE_ROLLBACK_STARTED | = 3 |

SOFTWARECOMPONENTACTION

Contains one record for each server Server Software Package component.

Table 11-7 *SoftwareComponentAction Field Names*

| Field Name | Type | Use |
|----------------------------|-----------|--|
| SOFTWARECOMPONENTACTIONKEY | integer | not null Assigned automatically: Default Auto increment. |
| PACKAGEACTIONID | integer | not null Links to the PACKAGEACTION table. |
| NAME | char | not null Name of the software component. |
| CREATIONDATE | timestamp | Time stamp of the action. |
| DESCRIPTION | varchar | The first record for the component the description is the description provided by the user when the component was created. As the components finish the description is one of those defined under "Valid Entries for DESCRIPTION" on page 376. |
| CODE | integer | Code representing the results of the action. For definitions, see "Valid Entries for CODE" on page 377. |
| ACTIONCODE | integer | Code representing the action being performed. For definitions, see "Valid Entries for ACTIONCODE" on page 377. |

Primary key (SOFTWARECOMPONENTACTIONKEY)

Valid Entries for DESCRIPTION

Did not meet requirements
Error processing requirements
Pre-install load/unload
Error pre-install load/unload
Pre-install scripts
Error pre-install scripts
Copy file changes
Error processing copy file
Text file changes
Error processing text files
NetWare SET parameters
Error processing NetWare SET parameters
Registry process
Error processing Registry
NetWare products process

Error in NetWare products process
Post-install script process
Error in post-install script process
Post-install load/unload process
Error in post-install load/unload process

Valid Entries for CODE

| | |
|---------|-----|
| Success | = 0 |
| Failure | = 1 |
| Partial | = 2 |

Valid Entries for ACTIONCODE

| | |
|-------------------|-------|
| Started | = 200 |
| Pre-Load | = 201 |
| Pre-Scripts | = 202 |
| Copy File Changes | = 203 |
| Text File Changes | = 204 |
| Set Parameters | = 205 |
| Registry | = 206 |
| Products.dat | = 207 |
| Post Scripts | = 208 |
| Post Load | = 209 |
| Requirements | = 210 |

Foreign Keys

Foreign keys set up relationships between tables.

POLICYACTION

"add foreign key (POLICYID) references POLICIES (POLICYID)"

POLICYACTION

"add foreign key (SERVERID) references SERVERS (SERVERID)"

PACKAGEACTION

"add foreign key (PACKAGEGUID) references PACKAGES (PACKAGEGUID)"

PACKAGEACTION

"add foreign key (SERVERID) references SERVERS (SERVERID)"

SOFTWARECOMPONENTACTION

"add foreign key (PACKAGEACTIONID) references PACKAGEACTION (PACKAGEACTIONID)"

11.4.3 Tiered Electronic Distribution Database Contents

Following are the database table definitions for Tiered Electronic Distribution:

- ◆ “TAB_NODE” on page 378
- ◆ “TAB_CHANNEL” on page 378
- ◆ “TAB_DISTRIBUTION” on page 379
- ◆ “TAB_DIST_VERSION” on page 379
- ◆ “TAB_DIST_ACTION” on page 380
- ◆ “TAB_CHANNEL_DISTRIBUTION” on page 381
- ◆ “Foreign Keys” on page 381

TAB_NODE

Contains one record for each Distributor, Subscriber, and External Subscriber in the tree.

Table 11-8 Tab_Node Field Names

| Field Name | Type | Use |
|-----------------|--------------------------------|---|
| ID | numeric(8,0) identity not null | Unique number automatically assigned. |
| NAME | varchar(255) not null | Tiered Electronic Distribution object DN. |
| TYPE | char not null | "D"=Distributor "T"=Subscriber (Transceiver) |
| NETWORK_ADDRESS | varchar(255) | IP address of server. |
| SERVER_NAME | varchar(255) | Not currently used. |

Primary key (ID)

Unique (NAME)

TAB_CHANNEL

Contains one record for each Channel object in the tree.

Table 11-9 Tab_Channel Field Names

| Field Name | Type | Use |
|------------|--------------------------------|---------------------------------------|
| ID | numeric(8,0) identity not null | Unique number automatically assigned. |

| Field Name | Type | Use |
|------------|--------------|--------------------------------|
| NAME | varchar(255) | not null DN of Channel object. |

Primary key (ID)
Unique (NAME)

TAB_DISTRIBUTION

Contains one record for each Distribution object in eDirectory.

Table 11-10 *Tab_Distribution Field Names*

| Field Name | Type | Use |
|----------------|-----------------------|--|
| ID | numeric(8,0) identity | not null Unique number automatically assigned. |
| NAME | varchar(255) | not null DN of Distribution object. |
| DISTRIBUTOR_ID | numeric(8,0) | not null Links to the TAB_NODE table. |

Primary key (ID)
Unique (NAME)

TAB_DIST_VERSION

Contains one record for each version of a Distribution and it is linked to the TAB_DISTRIBUTION table.

Table 11-11 *Tab_Dist_Version Field Names*

| Field Name | Type | Use |
|-----------------|------------------------|--|
| ID | numeric(10,0) identity | not null Unique number automatically assigned. |
| DISTRIBUTION_ID | numeric(8,0) | not null Links to the TAB_DISTRIBUTION table. |
| VERSION | bigint | not null Time stamp of the version. |
| SIZE | integer | not null Size of <code>distfile.ted</code> (the file containing the Distribution). |
| TIMESTAMP | datetime | not null Time stamp when the entry was made to the database. |
| DIRECT_ROUTING | bit | not null Not used at the current time. |
| LATEST_VERSION | bit | not null Latest version of this Distribution. Used internally to keep track of the latest version. |

Primary key (ID)
 Unique (DISTRIBUTION_ID, VERSION)

TAB_DIST_ACTION

Contains multiple records for each Distribution version for Send, Received, and Extracted.

Table 11-12 *Tab_Dist_Action Field Names*

| Field Name | Type | Use |
|------------------|----------------|---|
| ID | numeric(12, 0) | identity not null Unique number automatically assigned. |
| DIST_VERSION_ID | numeric(10, 0) | not null Links to the TAB_DIST_VERSION table. |
| NODE_ID | numeric(8,0) | not null Links to the TAB_NODE table for the node performing the following tasks: Create Send Receive Extract Post process |
| TIMESTAMP | datetime | not null Time stamp when the action was logged into the database. |
| STAGE | char | not null "C"=Create "S"=Send "R"=Receive "E"=Extract "P"=Post process |
| STATUS | char | not null "S"=Success "F"=Failure "P"=In process |
| STATUS_TIMESTAMP | datetime | not null Time stamp when the record was updated. |
| REASON_TEXT | varchar(255) | Reason for success or failure. For definitions, see "Valid Entries for REASON_TEXT" on page 380. |
| CHANEL_DIST_ID | numeric(8,0) | Links to the TAB_CHANNEL_DISTRIBUTION table. |

Primary key (ID)

Valid Entries for REASON_TEXT

The following are valid entries for the REASON_TEXT field name:

- ◆ "The Distribution was not received because this Subscriber does not meet the platform restrictions."

Self-explanatory.

- ◆ “The Distribution was shut down before it was received.”

This one is received in one of two situations: 1) there is a new configuration on the Subscriber so it needs to be updated before it can receive the Distribution; or, 2) there is a signature exception, such as the Subscriber cannot trust the Distribution came from a Distributor it trusts.

- ◆ “The Distribution was terminated before it was received.”

The Distribution was cancelled for a controlled reason.

- ◆ “There was an error receiving the Distribution.”

Something unexpected failed. For example, a socket exception, transport exception, and so on.

TAB_CHANNEL_DISTRIBUTION

Contains one record for each Channel/Distribution.

Table 11-13 *Tab_Channel_Distribution Field Names*

| Field Name | Type | Use |
|-----------------|--------------------------------|--|
| ID | numeric(8,0) identity not null | Unique number automatically assigned |
| CHANNEL_ID | numeric(8,0) | not null Links to the TAB_CHANNEL table. |
| DISTRIBUTION_ID | numeric(8,0) | not null Links to the TAB_DISTRIBUTION table. |
| TIMESTAMP | datetime | not null Time stamp for when the Distribution was built. |

Primary key (ID)

Unique (CHANNEL_ID, DISTRIBUTION_ID)

Foreign Keys

Foreign keys set up relationships between tables.

TAB_DISTRIBUTION

```
" add foreign key FK_TAB_DIST_REF_591_TAB_NODE (DISTRIBUTOR_ID)" + " references  
TAB_NODE (ID) on update restrict on delete restrict;"
```

TAB_DIST_VERSION

```
" add foreign key FK_TAB_DIST_REF_37_TAB_NODE (DISTRIBUTOR_ID)" + " references  
TAB_DISTRIBUTION (ID) on update restrict on delete restrict;"
```

TAB_DIST_ACTION

```
" add foreign key FK_TAB_DIST_REF_380_TAB_NODE (DIST_VERSION_ID)" + " references  
TAB_DIST_VERSION (ID) on update restrict on delete restrict;"
```

TAB_DIST_ACTION

```
" add foreign key FK_TAB_DIST_REF_1525_TAB_NODE (NODE_ID)" + " references  
TAB_NODE (ID) on update restrict on delete restrict;"
```

TAB_CHANNEL_DISTRIBUTION

```
" add foreign key FK_TAB_DIST_REF_572_TAB_DIST (DISTRIBUTION_ID)" + " references  
TAB_DISTRIBUTION (ID) on update restrict on delete restrict;"
```

TAB_CHANNEL_DISTRIBUTION

```
" add foreign key FK_TAB_DIST_REF_572_TAB_CHAN (CHANNEL_ID)" + " references  
TAB_CHANNEL (ID) on update restrict on delete restrict;"
```

Distribution Types

A

The following sections describe how to configure the Distribution types:

- ◆ [Section A.1, “Desktop Application,” on page 383](#)
- ◆ [Section A.2, “File,” on page 383](#)
- ◆ [Section A.3, “FTP,” on page 387](#)
- ◆ [Section A.4, “HTTP,” on page 390](#)
- ◆ [Section A.5, “MSI,” on page 392](#)
- ◆ [Section A.6, “Policy Package,” on page 395](#)
- ◆ [Section A.7, “RPM,” on page 396](#)
- ◆ [Section A.8, “Software Package,” on page 397](#)

A.1 Desktop Application

Use this option when the Distribution consists of an application created in ZENworks Desktop Management.

To create a Desktop Application Distribution:

- 1 Click the *Setup* button.

The Desktop Application Distribution Wizard is started.

After running the wizard the first time to create the Desktop Application Distribution, the *Setup* button is renamed to *Modify*.

For information on using the wizard, see [Step 7](#) under [Section 6.3, “Creating a Desktop Application Distribution,” on page 289](#).

After you exit the wizard, the *Current Configuration* field displays the current configuration of the Desktop Application Distribution. This is same information that is displayed on the Summary page of the Desktop Application Distribution Wizard.

- 2 To modify the Desktop Application Distribution’s configuration, click *Modify*.

This opens the Desktop Application Distribution Wizard again, where you can change the displayed configuration.

A.2 File

This option distributes files from the Distributor’s file system. Files cannot be gathered from locations accessed by way of mapped drives or UNC paths. Files from other servers can be distributed using the [FTP](#), [HTTP](#), and [RPM](#) types of Distribution.

With this type you can select files and directories for the Distribution and select a destination path for extraction on the Subscriber.

Use the following fields and buttons to configure a File Distribution:

- ◆ [Section A.2.1, “Files to Be Distributed,” on page 384](#)
- ◆ [Section A.2.2, “New Target,” on page 384](#)
- ◆ [Section A.2.3, “Add Directory,” on page 385](#)
- ◆ [Section A.2.4, “Add Files,” on page 385](#)
- ◆ [Section A.2.5, “Delete,” on page 385](#)
- ◆ [Section A.2.6, “Synchronize/Desynchronize,” on page 386](#)
- ◆ [Section A.2.7, “Verify Distributions,” on page 386](#)
- ◆ [Section A.2.8, “Maintain Trustees,” on page 387](#)
- ◆ [Section A.2.9, “Extract Error Handling,” on page 387](#)

A.2.1 Files to Be Distributed

An expandable tree structure showing target paths to the files to be distributed.

Modify the Distribution’s content as needed, then either click *Apply* to save the updated Distribution contents list, or click *OK* to save the changes and exit the Distribution’s properties.

A.2.2 New Target

If you do not want to use the default variable `%DEST_VOLUME%`, replace it with a target location indicating where you want the files to be distributed; for example:

- ◆ A volume on a NetWare® server, such as `data:\`
- ◆ A drive on a Windows server, such as `D:\`
- ◆ A file system on a Linux or Solaris server, such as `/`

Otherwise, press *Enter* to accept the variable. You can also replace it with a different variable.

IMPORTANT: If you use a UNC path, all Distributions are sent to only that one location. Instead, use variables. For more information, see [Chapter 9, “Variables,” on page 341](#).

The target can also be a full or partial path, not just the root designation of volume, drive, or root. For example:

- ◆ A volume on a NetWare server, such as `data:\file_distribution\files`
- ◆ A drive on a Windows server, such as `d:\file_distribution\files`
- ◆ A shared folder on a Windows server, such as `\\myserver\files`
- ◆ A file system on a Linux or Solaris server, such as `/user/file_distribution/files`

You must press *Enter* to add the target location change to the Distribution contents list.

Any target entries that do not exist in the target location are automatically created by the Distribution.

A.2.3 Add Directory

This option builds the target path. If you did not provide the full path using the *New Target* option, replace *New Directory* with a directory name.

IMPORTANT: If a directory has a % character as part of its name, you must enter two consecutive % characters (%%) so that the second character is recognized as a literal % character and not a variable indicator.

You can use this option to add multiple directories to create the full path. You can also add sibling directories to create additional paths, as determined by the subdirectory you select before clicking *Add Directory*.

You must press *Enter* each time to add the target directory change to the path.

A.2.4 Add Files

Use this option to browse for directories or files on the Distributor's file system that you want copied to the target Subscriber's file system.

Each directory or file you select is displayed with the full path on the source file system. This path identifies where to obtain the directory or file for copying to the target file system. The only path that is created on the target file system is the one you create using the *New Target* and *Add Directory* buttons, including any directories that you select with the *Add Files* button.

If you select a directory, all files and subdirectories under it are also selected for copying. Unlike the Copy File component in the Server Software Package, you cannot prune files and subdirectories from a selected directory. Any directory you browse for and add is not expandable in this view.

You can create multiple paths (sibling directories) at any point in a particular path, but you can only have one root location.

IMPORTANT: The directories that you select for the Distribution and any target directories cannot be Read-Only. File-writing to or from such directories will fail.

When you are finished adding files and target locations, either click *Apply* to save the updated Distribution contents list, or click *OK* to save the changes and exit the Distribution's properties.

A.2.5 Delete

Deletes whatever you have selected:

- ◆ **Root Location:** Removes all directories and files below it from the tree.
- ◆ **Directory:** Removes the directory and any of its files and subdirectories from the tree.
- ◆ **File:** Removes the file from the tree, but not from its hard disk location on the server.

Click *Apply* to save your changes, or click *OK* to save the changes and exit the Distribution's properties.

A.2.6 Synchronize/Desynchronize

This option causes the directories on the target server to be synchronized with the directories contained in the Distribution, or you can desynchronize them. When you specify a directory to be synchronized, it includes all files and subdirectories under the synchronized directory.

Directory synchronization does not affect adding of directories and files through a File Distribution. Synchronization's purpose is to delete files and directories on the Subscriber servers. In other words, if the directories and files on the target Subscriber server are not contained in the File Distribution, they are deleted if synchronization is enabled.

WARNING: If the target server contains directories not contained in the Distribution, those directories, subdirectories, and files are deleted from the target server's file system when the Distribution is extracted.

This can be very destructive, especially if the target directory is a root directory. Enable directory synchronization only where you are certain you want to allow existing directories that are not contained in the Distribution to be deleted.

Also, if the Distributor whose file system you are using for this Distribution is also a Subscriber that is subscribed to the Distribution, the Distributor's file system is treated like the other target Subscribers' file systems and can be deleted.

For more information on synchronization, see [Section 3.11.1, "Directory Sync Granularity for File Distributions,"](#) on page 176.

Directory synchronization provides granularity, where you can specify synchronization at any directory level in the Distribution to provide synchronization "from here down."

To synchronize directories, click either the target directory or the Distributor's source entry, then click *Synchronize*. A synchronization icon is added before the entry's graphic symbol.

The source location on the Distributor server determines what is kept in the synchronized directory on the Subscriber server, as depicted in the Distribution's *Files To Be Distributed* list. Thus, you are synchronizing the Subscriber's file system with the Distributor's, but only for the synchronized directory, with all its files and subdirectories.

The *Unsyncronize* button turns synchronization off (removing the icon), so that the specified directory is no longer synchronized.

If you enable synchronization, make sure that you refresh the Distributor before you make any changes to the directory being synchronized. This allows the Distributor to recognize that synchronization has been turned on, so that it rebuilds the Distribution with synchronization enabled.

A.2.7 Verify Distributions

Each time a Distribution changes, such as when files are modified or added, a new version is built and subsequently sent to the Subscribers. However, Subscribers might need to verify that the files contained in a Distribution have been extracted and installed to all Subscribers, even when there is no new version to send.

The verification option allows you to specify that if there is no new version of the Distribution to send, when the Send schedule starts the Distributor should send a request for the Subscriber to re-extract the current version to ensure that the files are installed.

A.2.8 Maintain Trustees

This option maintains each file's trustee attributes for the target NetWare file system so that they are the same as the source file system. The trustee information is obtained when the Distribution is built.

This is additive, meaning that it does not remove trustees on the target file system.

If synchronization is enabled for directories in a Distribution, the trustees of those directories are also synchronized.

A.2.9 Extract Error Handling

You have four options:

- ♦ **Fail on error:** Extraction of the Distribution stops. This results in a partial distribution. Correct the error and resend the Distribution.
- ♦ **Continue on error:** The extraction continues with an error written to the Subscriber's log file concerning the part of the extraction that failed.

By default, the Subscriber continues past error conditions, so as many files as possible are successfully extracted. You can avoid locked open file errors by selecting *Kill connection on open files*.

- ♦ **Retry ___ times:** By default, open files on the Subscriber server are not overwritten when the Distribution is extracted because the open files are locked by the operating system. As result, updated files in the Distribution are not replaced on the Subscriber server if they are open when the Distribution is extracted.

If you want the Subscriber to try multiple times to overwrite an open, locked file during extraction of the Distribution, specify the number of times the Subscriber should check the open file before failing to replace it because it is locked.

- ♦ **Kill connection on open files:** (NetWare only) Kills the connection that is holding the file open so that the file can be overwritten and the extraction can continue. (This applies only to files on the Subscriber server during extraction, not to files being accessed to build the Distribution.)

Server and NLM™ connections cannot be killed.

Error messages are written to the Subscriber log file.

A.3 FTP

This option distributes files from one or more FTP sources. Each source can contain one or more directories and/or files.

The FTP Distribution type enables the files to pass through your firewall as they are gathered into the Distribution.

If a target file is locked during extraction, the Subscriber throws an exception stating that the file could not be copied. The Distributor receives this information from the Subscriber and log the failure in the reporting database.

If you want to distribute an RPM software package from a Linux or Solaris FTP site, use the **RPM** type of Distribution rather than the FTP type.

Use the following fields and buttons to configure an FTP Distribution:

- ◆ [Section A.3.1, “Files To Be Distributed,” on page 388](#)
- ◆ [Section A.3.2, “New FTP Source,” on page 388](#)
- ◆ [Section A.3.3, “New Target,” on page 388](#)
- ◆ [Section A.3.4, “Add Directory,” on page 389](#)
- ◆ [Section A.3.5, “Add Files,” on page 389](#)
- ◆ [Section A.3.6, “Delete,” on page 389](#)
- ◆ [Section A.3.7, “Properties,” on page 390](#)
- ◆ [Section A.3.8, “Binary Transfer,” on page 390](#)
- ◆ [Section A.3.9, “Include Symbolic Link Files,” on page 390](#)

A.3.1 Files To Be Distributed

An expandable tree structure showing target paths to the FTP files to be distributed.

Modify the Distribution’s content as needed, then either click *Apply* to save the updated Distribution contents list, or click *OK* to save the changes and exit the Distribution’s properties.

A.3.2 New FTP Source

Specifies an FTP server from which to gather files for distribution:

- ◆ **FTP Server:** Specify the fully qualified host name of the FTP server.
- ◆ **Login Name:** Specify the login name that the Distributor Agent should use to access the FTP server. The default is Anonymous, which is often sufficient.
- ◆ **Password:** Provide the password for the login name. (You might need to scroll to the right to see the field.) For Anonymous, the password is typically your e-mail address.

A.3.3 New Target

If you do not want to use the default variable `%DEST_VOLUME%`, replace it with a target location indicating where you want the FTP files to be distributed; for example:

- ◆ A volume on a NetWare server, such as `data :` \
- ◆ A drive on a Windows server, such as `D :` \
- ◆ A file system on a Linux or Solaris server, such as `/`

Otherwise, press *Enter* to accept the variable. You can also replace it with a different variable.

The target can also be a full or partial path, not just the root designation of volume, drive, or root. For example:

- ◆ A volume on a NetWare server, such as `data:\ftp_distribution\files`
- ◆ A drive on a Windows server, such as `d:\ftp_distribution\files`
- ◆ A file system on a Linux or Solaris server, such as `/user/ftp_distribution/files`

You must press *Enter* to add the target location change to the Distribution contents list.

Any target entries that do not exist in the target location are automatically created by the Distribution.

A.3.4 Add Directory

This option builds the target path. If you did not provide the full path using the *New Target* option, replace *New Directory* with a directory name.

You can use this option to add multiple directories to create the full path. You can also add sibling directories to create additional paths, as determined by the subdirectory you select before clicking *Add Directory*.

You must press *Enter* each time to add the target directory change to the path.

A.3.5 Add Files

If you specified the correct FTP server information under *New FTP Source*, then you are provided access to the FTP server. Browse for the files. You can add multiple files.

If you added sibling directories using *Add Directory*, be sure to select those target paths and click *Add Files* to browse for the files to be distributed to those path locations.

When you are finished adding files and target locations, either click *Apply* to save the updated Distribution contents list, or click *OK* to save the changes and exit the Distribution's properties.

A.3.6 Delete

Deletes whatever you have selected:

- ◆ **Root Location:** Removes all directories and FTP files below it from the tree.
- ◆ **Directory:** Removes the directory and any of its FTP files and subdirectories from the tree.
- ◆ **File:** Removes the FTP file from the tree, but not from the FTP location.

Deleting any part of a target location, including a directory, means that the designated FTP files that were in the Distribution are no longer sent to the deleted location.

However, deleting FTP files from the Distribution means that the corresponding files are deleted from each Subscriber server when the Distribution is sent again and processed.

Click *Apply* to save your changes.

A.3.7 Properties

Displays the properties of the selected FTP source.

A.3.8 Binary Transfer

Enables file transfers in binary for when you are distributing executable files. Text files do not require a binary transfer.

A.3.9 Include Symbolic Link Files

If you want all symbolic link files in the added directory (including from all of its subdirectories) to be part of the FTP Distribution, select this box. This applies only when you add a directory.

You do not need to check this box for individually added symbolic link files.

This check box does not have any control over whether symbolic link files are displayed when browsing to add directories and files.

Click *Apply* to save your changes, or click *OK* to save the changes and exit the Distribution's properties.

A.4 HTTP

This option distributes files from one or more HTTP sources. Each source can contain one or more target entries.

The HTTP Distribution type enables the files to pass through your firewall as they are gathered into the Distribution.

If a target file is locked during extraction, the Subscriber throws an exception stating that the file could not be copied. The Distributor receives this information from the Subscriber and logs the failure in the reporting database.

Use the following fields and buttons to configure the Distribution:

- ◆ [Section A.4.1, "Files To Be Distributed," on page 390](#)
- ◆ [Section A.4.2, "New Target," on page 391](#)
- ◆ [Section A.4.3, "Add Directory," on page 391](#)
- ◆ [Section A.4.4, "Add Files," on page 391](#)
- ◆ [Section A.4.5, "Delete," on page 391](#)

A.4.1 Files To Be Distributed

An expandable tree structure showing target paths and the URLs to the HTTP files to be distributed.

Modify the Distribution's content as needed, then either click *Apply* to save the updated Distribution contents list, or click *OK* to save the changes and exit the Distribution's properties.

A.4.2 New Target

If you do not want to use the default variable `%DEST_VOLUME%`, replace it with a target location indicating where you want the HTTP files to be distributed; for example:

- ◆ A volume on a NetWare server, such as `data :` \
- ◆ A drive on a Windows server, such as `D :` \
- ◆ A file system on a Linux or Solaris server, such as `/`

Otherwise, press *Enter* to accept the variable. You can also replace it with a different variable.

The target can also be a full or partial path, not just the root designation of volume, drive, or root. For example:

- ◆ A volume on a NetWare server, such as `data : \http_distribution\files`
- ◆ A drive on a Windows server, such as `d : \http_distribution\files`
- ◆ A file system on a Linux or Solaris server, such as `/user/http_distribution/files`

You must press *Enter* to add the target location change to the Distribution contents list.

Any target entries that do not exist in the target location are automatically created by the Distribution.

A.4.3 Add Directory

This option builds the target path. If you did not provide the full path using the *New Target* option, replace *New Directory* with a directory name.

You can use this option to add multiple directories with this option to create the full path. You can also add sibling directories to create additional paths, as determined by the subdirectory you select before clicking *Add Directory*.

You must press *Enter* each time to add the target directory change to the path.

A.4.4 Add Files

Specify the URL where the files are available. You can specify a URL to a directory where multiple files exist, or include the filename in the URL to distribute a specific file.

To add files, you can select any directory in the path to assign the URL to it, including the root designation (such as `data :` \, `D :` \, or `/`).

If you added sibling directories using *Add Directory*, be sure to select those target paths and click *Add Files* to specify the URL for the files to be distributed to those path locations.

When you are finished adding files and target locations, either click *Apply* to save the updated Distribution contents list, or click *OK* to save the changes and exit the Distribution's properties.

A.4.5 Delete

Deletes whatever you have selected:

- ◆ **Root Location:** Removes all directories and files below it from the tree.

- ◆ **Directory:** Removes the directory and any of its URLs and subdirectories from the tree.
- ◆ **File:** Removes the URL from the tree, but not the files from the HTTP location.

Deleting any part of a target location, including a directory, means that the files designated by the URL that were in the Distribution are no longer sent to the deleted location.

However, deleting URLs from the Distribution means that the corresponding files are deleted from each Subscriber server when the Distribution is sent again and processed.

Click *Apply* to save your changes, or click *OK* to save the changes and exit the Distribution's properties.

A.5 MSI

This option distributes Microsoft Software Installer (MSI) packages to Windows servers for any Windows-based application, where the MSI engine is used to install the Windows software included in an MSI Distribution. You can create and configure MSI Distributions in both ConsoleOne and iManager.

Use the following to configure the Distribution:

- ◆ [Section A.5.1, "Adding," on page 392](#)
- ◆ [Section A.5.2, "Removing," on page 392](#)
- ◆ [Section A.5.3, "Configuring," on page 392](#)
- ◆ [Section A.5.4, "Rearranging," on page 395](#)

A.5.1 Adding

Add one or more MSI or MSP packages using one of the following:

Add from Distributor: The `.msi` and `.msp` files must reside on the file system of the Distributor server that owns this MSI Distribution.

Add from FTP site: The `.msi` and `.msp` files can be retrieved from an FTP site.

A.5.2 Removing

Select an MSI package from the list and click Remove to delete it from the list.

A.5.3 Configuring

To configure each MSI package, select the package in the *Selected Packages* column, click *Edit Parameter List* to open the Edit Parameters dialog box, then fill in the fields:

- ◆ ["Distribution Includes Box" on page 393](#)
- ◆ ["Options Box" on page 393](#)
- ◆ ["Transforms Box" on page 394](#)
- ◆ ["Custom Parameters Field" on page 394](#)
- ◆ ["Command Field" on page 394](#)

Distribution Includes Box

Select one of the following options:

Package file only: Include only the MSI package in the Distribution.

Package files and folders: Include the MSI package, all files located in the same folder, and all subfolders and files. This assumes that all of the necessary supporting files for an MSI package are included in its folder and subfolders.

Options Box

Select from the following options:

Install: Causes the MSI package to be installed.

Uninstall: Causes the MSI package to be uninstalled.

Patch: This field is dimmed because it applies only to an MSP package.

Administrative install: Causes the MSI package to be installed without deleting the MSI package (as standard practice), so that it can be available for a self-repair. This option is used in conjunction with an administrative image of the package. For more information, see the [InstallShield Tip from AdminStudio \(http://www.installshield.com/news/newsletter/0302-articles/setupexe.asp\)](http://www.installshield.com/news/newsletter/0302-articles/setupexe.asp).

Repair: If you select this option, select from the following check boxes. They are the common MSI flags that can be passed to the MSI engine to specify the types of repairs to be made:

- ◆ **Missing file:** Instructs Windows Installer to reinstall a file only if it is missing.
- ◆ **Older file:** Instructs Windows Installer to reinstall a file if it is missing or if the installed file's version is older than the file in the MSI package.
- ◆ **Equal or older file:** Instructs Windows Installer to reinstall a file if it is missing or if the installed file's version is the same as or older than the file in the MSI package.
- ◆ **Force all:** Instructs Windows Installer to reinstall all files.
- ◆ **Use registry keys:** Instructs Windows Installer to rewrite all per-user entries from the MSI package to the Windows system registry. Per-user entries are those entries contained in the HKEY_CURRENT_USER and HKEY_USERS registry hives.
- ◆ **Computer registry keys:** Instructs Windows Installer to rewrite all per-machine entries from the MSI package to the Windows system registry. Per-machine entries are those entries contained in the HKEY_LOCAL_MACHINE and HKEY_CLASSES_ROOT registry hives.
- ◆ **Failed checksum:** Instructs Windows Installer to perform a checksum on all executable files and to reinstall a file if it is missing or if the checksum verifies that the file is corrupt. Only files that have msidbFileAttributesChecksum in the Attributes column of the MSI package's File Table are repaired.
- ◆ **Install and re-cache:** Instructs Windows Installer to install files from the re-cache (local) source rather than the source package.
- ◆ **Shortcuts:** Instructs Windows Installer to reinstall the MSI application's shortcuts, overwriting any existing shortcuts and icons.
- ◆ **Different file version:** Instructs Windows Installer to reinstall a file if it is missing or if the installed file's version is not exactly the same as the file in the MSI package.

Transforms Box

To include transforms in the Distribution, select an MSI package (not an MSP package) to be transformed, select the Edit Parameters button, then do the following. Repeat these steps as necessary for each transform in an MSI package, and for each MSI package.

- ◆ [“Adding” on page 394](#)
- ◆ [“Removing” on page 394](#)
- ◆ [“Rearranging” on page 394](#)

Adding

Add one or more transforms using one of the following:

Add from Distributor: The `.mst` file must reside on the file system of the Distributor server that owns this MSI Distribution.

Add from FTP Site: The `.mst` file can be retrieved from an FTP site.

Transform files are used to modify the behavior of the MSI package that you selected in the Selected Packages column of the Type tab.

When two or more transforms are applied to the same MSI package property, it retains the value applied by the transform that was last applied.

For more information about creating and configuring transforms, see the documentation you received with the software application.

Removing

Select a transform from the list and click Remove to delete it from the list.

Rearranging

Use the Up and Down buttons to rearrange the order in which the transforms are applied.

When you rearrange the execution order, remember that an MSP package patches a specific MSI package, so it should be listed after the MSI package.

Custom Parameters Field

You can modify the listed command line parameters for the MSI package.

Some MSI Distributions can fail to extract on Windows 2000 servers. To solve this problem, see [“MSI Distribution Extraction Errors” on page 116](#).

Command Field

This field is display-only.

The parameters listed in this field are for the default options when you first view the Parameters dialog box. These parameters are automatically updated as you modify any options in the Distribution Includes, Options, or Transforms boxes, or add any parameters in the Custom Parameters field.

A.5.4 Rearranging

Use the Up and Down buttons to rearrange the order in which the MSI packages are applied.

A.6 Policy Package

Use this option when the Distribution consists of one or more policy packages containing enabled and configured policies. This is how Subscribers receive policies.

Only the policies contained in the Distributed Policy Package are distributed for enforcement on Subscriber servers. The Container Package and Service Location Package are not distributed; they continue to be associated for enforcement on Distributor servers.

To send a Policy Package Distribution to a Subscriber using an External Subscriber object, you must edit the `agentinfo.properties` file to prevent trusted tree errors. For more information, see [“Preventing Trusted Tree Errors for Policy Package Distributions” on page 161](#).

For information on creating specific policies, see [Chapter 4, “Server Policies,” on page 193](#).

Use the following fields and buttons to configure the Distribution:

- ◆ [Section A.6.1, “Up/Down,” on page 395](#)
- ◆ [Section A.6.2, “Add,” on page 395](#)
- ◆ [Section A.6.3, “Delete,” on page 395](#)
- ◆ [Section A.6.4, “Properties,” on page 395](#)
- ◆ [Section A.6.5, “The Following Policy Packages Will Be Distributed,” on page 395](#)

A.6.1 Up/Down

Rearranges the installation order for the policy packages.

A.6.2 Add

Adds a policy package to the Distribution.

A.6.3 Delete

Deletes the policy package from those listed.

A.6.4 Properties

Displays the properties of the selected policy package, which you can then edit.

A.6.5 The Following Policy Packages Will Be Distributed

Lists the policy packages to be distributed and the order of distribution.

A.7 RPM

You can distribute any Red Hat Package Manager (RPM) packages you have created to your Linux or Solaris servers through Tiered Electronic Distribution.

Use the following fields and buttons to configure the Distribution:

- ◆ [Section A.7.1, “Up/Down,” on page 396](#)
- ◆ [Section A.7.2, “Add From Distributor,” on page 396](#)
- ◆ [Section A.7.3, “Add From FTP Site,” on page 396](#)
- ◆ [Section A.7.4, “Delete,” on page 396](#)
- ◆ [Section A.7.5, “Selected Packages,” on page 396](#)
- ◆ [Section A.7.6, “Installation Parameters,” on page 396](#)

A.7.1 Up/Down

Arranges the installation order for the RPM packages.

A.7.2 Add From Distributor

Browse the Distributor’s file system and select the RPM packages.

A.7.3 Add From FTP Site

Browse the FTP site and select the RPM packages.

A.7.4 Delete

Deletes the selected RPM package from the list.

A.7.5 Selected Packages

Lists the RPM packages you have added.

A.7.6 Installation Parameters

Lists the RPM installation parameters you have added.

Important points when entering parameters:

- ◆ You must press Enter for parameter entries in the text field, or the entries are not saved.
- ◆ You cannot remove a single parameter once it has been entered; you must re-enter the entire parameter string without the one you wanted removed.
- ◆ You cannot change the case of a parameter and have that change recognized. Instead, change the parameter to a different character, then change it back again to the original character with the desired case.

A.8 Software Package

Use this option when the Distribution consists of one or more software packages created in the Server Software Package namespace in ConsoleOne.

Use the following fields and buttons to configure the Distribution:

- ◆ [Section A.8.1, “Up/Down,” on page 397](#)
- ◆ [Section A.8.2, “Add,” on page 397](#)
- ◆ [Section A.8.3, “Delete,” on page 397](#)
- ◆ [Section A.8.4, “Selected Software Packages,” on page 397](#)

A.8.1 Up/Down

Rearranges the installation order for the software packages.

A.8.2 Add

Adds a software package to the Distribution.

A.8.3 Delete

Deletes the software package from the list.

A.8.4 Selected Software Packages

Lists the software packages to be distributed and the order of distribution.

Schedule Types

B

Table B-1 describes each of the Novell® ZENworks® schedule types, with links to the steps for configuring them.

Table B-1 Schedule Types

| Schedule Type | Description |
|------------------|--|
| Daily | Runs the scheduled item daily. Daily includes specifying a run time window, running randomly within the window of time, and running repeatedly every xxx hours or minutes. Used by all Policy and Distribution Services components. |
| Event | Runs the scheduled policy according to the specified event, such as at system startup or shutdown, or a third-party application-defined event. Used only by policies. |
| Interval | Repeats running the scheduled item every xxx days, hours, minutes, and/or seconds. For Distributors only, the interval begins after the Distributor re-reads Novell eDirectory™. Any frequency from a few seconds to many days can be specified. Used by policies, Distributors, Distributions, Channels, and Subscribers. |
| Monthly | Runs the scheduled item on the selected day of the month. Monthly includes specifying a run time window and running randomly within the window of time. Used by all Policy and Distribution Services components. |
| Never | Prevents any of the four possible schedules from occurring. Only used by Tiered Electronic Distribution. This is generally used for manual control over a particular schedule. Typically, you do not need to leave an object configured with the Never schedule type for an extended period of time. If an object is no longer used, you can remove it using the Delete TED Object menu option in Novell ConsoleOne®. |
| Package Schedule | Runs the scheduled item according to the default schedule, which can be changed on the Policies tab. Used only by policies. |
| Relative | Runs the scheduled policy one time relative to a specified number of days, hours, minutes, and seconds from when the policy package is extracted. For example, if you set the time to one hour and refresh the Distributor, a new policy package is sent to the Subscriber, and it runs one hour after extraction. Used only by policies. Any time range, from a few seconds to many days, can be specified. |
| Run Immediately | Runs the scheduled item immediately upon refreshing the policy, beginning after the Distributor re-reads eDirectory. Includes repeating the action every xxx days, hours, minutes, and seconds. Any frequency from a few seconds to many days can be specified. Used only by policies, Distributions, Channels, and Subscribers. |
| Time | Runs the scheduled item once at the date and time specified. Used by all Policy and Distribution Services components. |
| Weekly | Runs the scheduled item on the selected day of the week. Weekly includes specifying a run time window, and running randomly within the window of time. Used only by policies. |

| Schedule Type | Description |
|---------------|---|
| Yearly | Runs the scheduled item on the selected day of the year. Yearly includes specifying a run time window, and running randomly within the window of time. Used by all Policy and Distribution Services components. |

B.1 Daily

To schedule an item to run daily:

- 1 Click the down-arrow on *Schedule Type* > select *Daily*, then select one or more days of the week.
- 2 In *Start Time*, select the schedule's starting time for the day.
- 3 In *End Time*, select the latest time in the day for the schedule to run.
- 4 To have the schedule start randomly during the selected time period, select the *Randomly Dispatch* check box.
- 5 To have the schedule repeat the action, select the check box for the *Repeat the Action Every* field, then select how often the action should be repeated.

You can leave any of the options zeroed, but you must have a value in at least one of the time increments.

- 6 Click *Apply* to save the change.

B.2 Event

To schedule a policy to run when an event happens:

- 1 Click the down-arrow on *Schedule Type*, select *Event*, then select the event to activate the schedule:

| Event | Description |
|-----------------|---|
| System Startup | Runs the action when the system starts up. |
| System Shutdown | Runs the action before the system shuts down. |
| Custom Event ID | Third-party application-defined event. |

- 2 Click *Apply* to save the change.

B.3 Interval

To schedule an item to run at an interval of time:

- 1 Click the down-arrow on *Schedule Type*, select *Interval*, then select the interval of time for repeating the action.

You can leave any of the options zeroed, but you must have a value in at least one of the time increments.

- 2 Click *Apply* to save the change.

B.4 Monthly

To schedule an item to run monthly:

- 1 Click the down-arrow on *Schedule Type*, select *Monthly*, select the option, then select the day of the month.

or

Select the option for the last day of the month (whether 28, 29, 30, or 31).
- 2 In *Start Time*, select the schedule's starting time for the day.
- 3 In *End Time*, select the latest time in the day for the schedule to run.
- 4 To have the schedule start randomly during the selected time period, select the *Randomly Dispatch* check box.
- 5 Click *Apply* to save the change.

B.5 Never

This is generally used for manual control over a particular schedule. Typically, you do not need to leave an object configured with the Never schedule type for an extended period of time. If an object is no longer used, you can remove it using the Delete TED Object menu option in ConsoleOne.

Never has the following effects on the distribution schedules:

- ♦ **Refresh (Distributor object):** Prevents the Distributor from reading eDirectory to discover new distribution work.
- ♦ **Build (Distribution object):** Prevents the Distributor from building that particular Distribution.
- ♦ **Send (Channel object):** Prevents all Distributions listed in the Channel from being sent.
- ♦ **Extract (Subscriber object):** Prevents the Subscriber from extracting any of the Distributions it has received but not yet extracted. However, the Subscriber server can still receive Distributions.

When Distributions are changed to another schedule, all Distributions not yet extracted are extracted by the Subscriber according to the new schedule. When temporarily overridden using the *ZENworks Server Management* role in Novell iManager, all Distributions not yet extracted by the Subscriber are then extracted.

In each of these cases, you can manually override the Never action in iManager (see [“Forcing Policy and Distribution Services Agent Actions” on page 79](#)). However, the Never type continues to be set for the schedule after that override action occurs.

To schedule a Tiered Electronic Distribution item to never run automatically:

- 1 Click the *Type* tab, then select the down-arrow on *Schedule Type*.
- 2 Select *Never*.
- 3 Click *Apply* to save the change.

B.6 Package Schedule

Each policy package has a default schedule for all policies in that package.

You do not need to do anything to schedule a policy to run according to the current Default Package Schedule.

To change the Package Schedule:

- 1 In ConsoleOne, select the OU containing your server policies, right-click the Distributed Server Package (in the right pane), then click *Properties*.
- 2 Click *Edit*.
- 3 Change *Package Schedule* to one of the following:

| | | |
|---------|---------------|----------|
| Daily | Yearly | Event |
| Weekly | Relative | Interval |
| Monthly | Run Immediate | Time |

- 4 Click *Apply* to save the change.

B.7 Relative

To schedule a policy to run relative to the time the policy package has been extracted:

- 1 Click the down-arrow on *Schedule Type*, select *Relative*, then select an amount of time.

You can leave any of the options zeroed, but you must have a value in at least one of the time increments.

- 2 Click *Apply* to save the change.

B.8 Run Immediately

To schedule an item to run immediately:

- 1 Click the down-arrow on *Schedule Type*, then select *Run Immediately*.
- 2 If you want to repeat the action, select the *Repeat* check box.
- 3 Select a length of time.

You can leave any of the options zeroed, but you must have a value in at least one of the time increments.

- 4 Click *Apply* to save the change.

B.9 Time

To schedule an item to run at a specific time:

- 1 Click the down-arrow on *Schedule Type*, select *Time*, then select the calendar icon.
- 2 In the *Select Date and Time* dialog box:
 - 2a Select the month.
 - 2b Select the year.
 - 2c Select the day of the month.
 - 2d Select the time of day, then click *OK*.

- 3 Click *Apply* to save the change.

B.10 Weekly

To schedule a policy to run weekly:

- 1 Click the down-arrow on *Schedule Type*, select *Weekly*, then select one day of the week.
- 2 In *Start Time*, select the schedule's starting time for the day.
- 3 In *End Time*, select the latest time in the day the schedule can run.
- 4 To have the schedule start randomly during the selected time period, select the *Randomly Dispatch* check box.
- 5 Click *Apply* to save the change.

B.11 Yearly

To schedule an item to run yearly:

- 1 Click the down-arrow on *Schedule Type*, select *Yearly*, then select the calendar icon.
- 2 In the Select Date dialog box:
 - 2a Select the month.
 - 2b Select the day of the month.
- 3 In *Start Time*, select the schedule's starting time for the day.
- 4 In *End Time*, select the latest time in the day the schedule can run.
- 5 To have the schedule start randomly during the selected time period, select the *Randomly Dispatch* check box.
- 6 Click *Apply* to save the change.

Server Console Commands

C

You can perform some of the Novell® ZENworks® Server Management functions using command line entries on a NetWare® server console. The server commands documented here are those that are applicable to Server Management's Server Policies and Tiered Electronic Distribution.

For ways to perform the server console commands in a Web browser using the ZENworks Server Management role in Novell iManager, see [Chapter 2, "Novell iManager," on page 63](#).

A ZENworks Server Management console command that is entered on a server console is executed only on that server. For more information, review the following sections:

- ♦ [Section C.1, "ZENworks Server Management Console Commands," on page 405](#)
- ♦ [Section C.2, "Java Console Commands," on page 408](#)

C.1 ZENworks Server Management Console Commands

[Table C-1](#) lists the ZENworks Server Management server console commands with short descriptions of the commands. The table also indicates at which server console prompt you can give a command.

The column heading M is for the server's main console prompt and Z for the ZENworks Server Management prompt. Under a console prompt column, a Y indicates that you can issue the command at that prompt and a – indicates that you cannot issue the command at that prompt.

Table C-1 Console Commands

| Command | M | Z | Description |
|---------------------|---|---|--|
| HELP | Y | Y | Displays a list of available commands. Only the commands applicable to a component are displayed. |
| HELP <i>command</i> | Y | Y | Displays help for the specified command. |
| CLS | Y | Y | Clears the screen. Useful for quickly recognizing which information is new when you enter a command. |

| Command | M | Z | Description |
|-----------------------|---|---|--|
| DOWN <i>option</i> | Y | Y | <p>This is similar to the command used on the server's main console prompt. However, if you use DOWN at the ZENworks Server Management prompt, server policy settings for downing the server are followed.</p> <p>For the ZENworks Server Management prompt, this command has several options:</p> <ul style="list-style-type: none"> ◆ DOWN SERVER: Downs the server only; does not bring it back up. ◆ DOWN STATUS: Displays the current down status. ◆ DOWN RESTART: Downs the server, then restarts it. ◆ DOWN RESET: Downs the server, then resets it. ◆ DOWN CANCEL: Allows you to cancel the down, up to when the server is actually taken down. This does not leave the server in an unusable state. ◆ DOWN !: Causes the down process to execute immediately, ignoring the Down Server Process policy that might be in effect. |
| EVENTS <i>option</i> | - | Y | <p>The command has three options:</p> <ul style="list-style-type: none"> ◆ EVENTS LIST: Lists all registered events, including third-party events. ◆ EVENTS STATUS: Gives the status of each event. ◆ EVENTS FIRE <i>event_ID</i>: Allows you to manually run an event. |
| EXIT | - | Y | <p>Closes the current command prompt's Java software. For example, if given at the Subscriber prompt, the Subscriber's Java software is closed.</p> |
| LISTPLUGINS | - | Y | <p>Lists the current Server Management plug-ins.</p> |
| PACKAGE <i>option</i> | - | Y | <p>You can do the following for the software packages installed on the server:</p> <ul style="list-style-type: none"> ◆ PACKAGE LIST: Lists the currently installed software packages. This is useful for knowing which packages to roll back and the order that they should be rolled back, which is the reverse order in which they finished installing, not the order that they started installing. ◆ PACKAGE PROCESS <i>full_package_path</i>: Use this to manually install a software package. ◆ PACKAGE ROLLBACK: Automatically rolls back (uninstalls) the most recently installed software package. For example, you installed three software packages on a server (Package1, Package2, and Package3), and Package1 was installed first, Package2 second, and Package3 last. If you want to roll back Package2, you need to first roll back Package3. To do so, enter <code>package rollback</code> at the server console once for Package3, then again for Package2. <p>The software package installation order is not guaranteed, because the order is determined by when a package has finished processing. Therefore, the installation order might be Package2, Package1, Package3 when using the Package Rollback command. This order is shown by the Package List command.</p> |

| Command | M | Z | Description |
|--|---|---|---|
| POLICY or POLICY LIST | - | Y | Lists the effective server policies. Each policy listed has a corresponding policy number for reference when using the POLICY ENFORCE command. |
| POLICY ENFORCE <i>policy_number</i> | - | Y | Used to manually enforce a specific policy. You can find the <i>policy_number</i> using the POLICY LIST command. This is useful for enforcing a policy ahead of its schedule. However, you usually use POLICY REFRESH first to ensure you are enforcing the most recent changes. |
| POLICY ENFORCE ALL | - | Y | Used to manually enforce all effective policies, such as after doing a POLICY REFRESH. |
| POLICY EVENTBASED | - | Y | Lists the event-based policies. |
| POLICY PLUGINS | - | Y | Lists the current policy enforcers and the current event handlers. |
| POLICY REFRESH | - | Y | Refreshes only the server's policies and schedules, as required (unlike the REFRESH command, which refreshes policies and undoes any changes made to the prompts). After using this command, you should do a POLICY ENFORCE. |
| POLICY REFRESHONLY | - | Y | Refreshes the server's policies, but does not schedule effective policies. |
| POLICY RESCHEDULEONLY | - | Y | Reschedules all current policies according to their schedules. Does not refresh the effective policies. |
| POLICY SCHEDULES | - | Y | Lists all policy schedules that are in effect. |
| PROMPT | - | Y | Temporarily resets the current prompt. It reverts back to whatever is specified in the Novell eDirectory™ object for the console prompt when the Java process is exited or restarted, or when the REFRESH command is given. |
| REFRESH | - | Y | Manually forces a refresh of a policy, including pending changes to service locations for the current server and temporary changes to ZENworks Server Management prompts. Used alone, it refreshes only the ZENworks Server Management policy. Use POLICY REFRESH to refresh all policies. |
| SETCONSOLELEVEL <i>number</i> | - | - | Sets the console message level: 0: No messages 1: Errors 2: Successes & level 1 messages 3: Warnings & level 2 messages 4: Information & level 3 messages 5: Trace information & level 4 messages 6: Developer trace information & level 5 messages |

| Command | M | Z | Description |
|-------------------------------|---|---|--|
| SETFILELEVEL <i>number</i> | - | Y | Sets the file message level: 0: No messages 1: Errors 2: Successes & level 1 messages 3: Warnings & level 2 messages 4: Information & level 3 messages 5: Trace information & level 4 messages 6: Developer trace information & level 5 messages |
| SHOWSCHEDULE | - | Y | Lists the current schedules. |
| SHOWVARS | - | Y | Lists the predefined variables and their values. These variables can be used in Server Software Packages. |
| STATUS | - | Y | Lists the current status of Policy and Distribution Services, including: Base Path Plug-ins Loaded Events Registered Scheduled Items Console Level |
| TIME | Y | Y | Returns the current date and time that the server is set to. |
| VERSION | Y | Y | Returns the Server Management version for the ZENworks Server Management prompt, and the NetWare version for the console's main prompt. |

C.2 Java Console Commands

Table C-2 lists some useful Java Virtual Machine (JVM*) commands.

Table C-2 Java Commands

| Command | Description |
|-----------------------|--|
| java -show | Lists all loaded Java processes. |
| java -kill <i>nnn</i> | Kills the specified Java process. (<i>nnn</i> represents the Java process number from the <code>java -show</code> listing.) |
| java -killzfsexit | Kills all Server Management Java processes. |
| java -killall | Stops all loaded Java processes; however, it leaves Java loaded. |
| java -version | Displays the JVM version. |

| Command | Description |
|---|--|
| <code>java -exit</code> or <code>unload java</code> | This attempts to unload all Java process, including the JVM. <code>Java -exit</code> is the preferred command. This command is required for unloading any native NLM™ files that are called from Java, such as <code>zenfile.nlm</code> . |

Load/Unload Actions

D

This information is used in several setup steps for the Server Policies components (see [Chapter 4, “Server Policies,”](#) on page 193) or Server Software Packages (see [Chapter 5, “Server Software Packages,”](#) on page 237).

- ◆ [Section D.1, “Load NLM/Process,”](#) on page 411
- ◆ [Section D.2, “Load Java Class,”](#) on page 411
- ◆ [Section D.3, “Unload Process,”](#) on page 412
- ◆ [Section D.4, “Start Service,”](#) on page 412
- ◆ [Section D.5, “Stop Service,”](#) on page 412

D.1 Load NLM/Process

For all supported platforms.

If you select an NLM™ to be loaded by the software package, and the NLM is already running on the target server, the package installation fails and is rolled back (if rollback is enabled).

You can make sure that an NLM is not already loaded when you are including it in the software package by adding an unload option for that NLM before adding the load option, but only if this NLM does not require user input from the keyboard to unload it.

Filename: This must be the exact name, including the full path to the executable, unless the path to the file is a system path variable. For NLM files, include the `.nlm` extension.

Parameters: Include any command line parameters for the NLM or process being run.

Wait for this process to terminate before continuing: You can select this option for an NLM or process that terminates itself. It must terminate within 10 minutes or the whole loading process fails. By default, this option is dimmed.

D.2 Load Java Class

For NetWare® only.

Filename: This must be the exact class name as listed in the JAR file’s source code. For example, `cpkmidtier.CPKMidTierConfigure` is used by the JAR file listed in the example shown below in the *JVM Parameters* field.

The `.class` extension is not necessary.

Parameters: Include any command line parameters required by the Java application being run, such as the Windows variable `%computername%`.

JVM parameters: Include any parameters for the `java` command. For example, the following parameters are the same as used with the `java` command on a Windows command prompt:

```
-cp -DSystemRoot=%SystemRoot% -DSystemDrive=%SystemDrive% -  
Djava.library.path=C:\onenet C:\onenet\CPKMidTierConfigure.jar
```

In this example, %systemdrive% is a Windows system variable and your JAR file is named CPKMidTierConfigure.jar. The class entered in the *Filename* field (such as cpkmidtier.CPKMidTierConfigure) is used by the CPKMidTierConfigure.jar file, and this JAR file needs a Djava library path to be C:\onenet. Copy your JAR file into the C:\onenet\ directory for the CPK to find it.

Wait for this process to terminate before continuing: You can select this option for a Java application that terminates itself. There is no time limit. It waits as long as the application is running. By default, this option is dimmed.

D.3 Unload Process

For all supported platforms.

If the NLM requires intervention to unload, you must remember to unload it manually before trying to install the software package.

Filename: This must be the exact name (the full path is not required). Because many NLM programs require user input to unload, their unloading cannot be automated.

Wait for this process to unload before continuing: You can select this option for a process that unloads itself. By default, this option is dimmed.

D.4 Start Service

For Windows only.

Service name: This must be the exact name.

Wait for this service to finish running before continuing: You can select this option for a service that starts itself. By default, this option is dimmed.

D.5 Stop Service

For Windows only.

Service name: This must be the exact name.

Wait for this service to stop before continuing: You can select this option for a service that stops itself. By default, this option is not selected.

Requirements for Server Software Packages

E

This information is used in several setup steps for software packages. For more information, see [Chapter 5, “Server Software Packages,” on page 237](#).

IMPORTANT: By selecting a requirement, you are prescribing that it must be met to allow the software package or package component to be installed.

[Table E-1](#) lists the requirement types:

Table E-1 *Server Software Package Requirements*

| Requirement | Description |
|----------------------------------|---|
| Operating System | The operating system (OS) requirements for running the files in the software package, including both the OS the files need for running and whether the target server has that OS. |
| Memory (RAM) | The minimum RAM required for running the files in the software package. If the target server does not meet that minimum, the software package is not distributed to it. |
| Disk Space | The minimum free disk space required for installing the files on the target server. If the target server does not meet that minimum free space, the software package is not distributed to it. |
| SET Commands | Which NetWare® SET commands you want specifically configured on the target server for the software package. |
| Registry | The registry changes that can be required on the target server for the files in the software package. For information on configuring individual registry entries, see Appendix F, “Registry Entries for Server Software Package Components,” on page 421 . |
| File | Indicates whether a file on the target server should exist or have a certain date. |
| Products.dat | Changes to <code>products.dat</code> that the software package requires. Usually, the changes are to update the versions of the software on the server from the contents of the software package. The <code>products.dat</code> file is used to determine which software and which version exist on the server. |

E.1 Operating System

You can require the server to have a certain operating system before installing the software package.

To configure the server operating system requirement:

- 1 With the operating system requirement selected, select the server’s platform.

Available platforms are NetWare, Windows, Linux, and Solaris.

2 Select the version relationship:

- Any
- Less Than
- Less Than or Equal To
- Equal To
- Greater Than
- Greater Than or Equal To

3 If you select an option other than *Any* for the *Version* field, fill in the *Major*, *Minor*, and *Revision* fields according to the information in the following table.

For Windows servers, version information cannot be specified. Therefore, Windows is not included in the table.

The *Major* and *Minor* fields are for the upper version limit. For Netware and Windows, the *Revision* field is for the required service pack revision. For Linux, the *Revision* field is for the Linux distribution update version.

| Operating System Version | Subscriber Version ¹ | Major | Minor | Revision |
|--------------------------|--|-------|-------|----------|
| NetWare 5.1 + SP5 | ZENworks 7, ZENworks 7 w/SP1, ZENworks 6.5, ZENworks 6.5 SP1, ZENworks 6.5 SP2, or ZfS 3.0.2 | 5 | 10 | 5 |
| NetWare 5.1 + SP6 | ZENworks 7, ZENworks 7 w/SP1, ZENworks 6.5, ZENworks 6.5 SP1, ZENworks 6.5 SP2, or ZfS 3.0.2 | 5 | 10 | 6 |
| NetWare 5.1 + SP7 | ZENworks 7, ZENworks 7 w/SP1, ZENworks 6.5, ZENworks 6.5 SP1, or ZENworks 6.5 SP2 | 5 | 10 | 7 |
| NetWare 5.1 + SP8 | ZENworks 7 or 7 w/SP1 only | 5 | 10 | 8 |
| NetWare 6 + SP2 | ZfS 3.0.2 only | 6 | 0 | 2 |
| NetWare 6 + SP3 | ZfS 3.0.2 only (with JVM 1.4.1) | 6 | 0 | 3 |
| NetWare 6 + SP4 | ZENworks 7, ZENworks 7 w/SP1, ZENworks 6.5, ZENworks 6.5 SP1, ZENworks 6.5 SP2, or ZfS 3.0.2 | 6 | 0 | 4 |
| NetWare 6 + SP5 | ZENworks 7, ZENworks 7 w/SP1, ZENworks 6.5, ZENworks 6.5 SP1, or ZENworks 6.5 SP2 | 6 | 0 | 5 |
| NetWare 6.5 | ZfS 3.0.2 only | 6 | 5 | 0 |
| NetWare 6.5 + SP1a | ZENworks 6.5 or ZfS 3.0.2 | 6 | 5 | 1 |
| NetWare 6.5 + SP1.1 | ZENworks 7, ZENworks 7 w/SP1, ZENworks 6.5, ZENworks 6.5 SP1, ZENworks 6.5 SP2, or ZfS 3.0.2 | 6 | 5 | 1 |
| NetWare 6.5 + SP2 | ZENworks 7, ZENworks 7 w/SP1, ZENworks 6.5, ZENworks 6.5 SP1, or ZENworks 6.5 SP2 | 6 | 5 | 2 |
| NetWare 6.5 + SP3 | ZENworks 7, ZENworks 7 w/SP1, ZENworks 6.5 SP1, or ZENworks 6.5 SP2 | 6 | 5 | 3 |
| NetWare 6.5 + SP4 | ZENworks 7 and 7 w/SP1 only | 6 | 5 | 4 |

| Operating System Version | Subscriber Version ¹ | Major | Minor | Revision |
|--|--|-------|-------|------------------------------|
| NetWare 6.5 + SP5 | ZENworks 7 w/SP1 only | 6 | 5 | 5 |
| OES NetWare | ZENworks 7 and 7 w/SP1 only | 6 | 5 | 4 |
| OES NetWare + SP1 | ZENworks 7 and 7 w/SP1 only | 6 | 5 | 5 |
| OES NetWare + SP2 | ZENworks 7 SP1 only | 6 | 5 | 6 |
| OES Linux | ZENworks 7 and 7 w/SP1 only | 2 | 6 | <i>variable</i> ² |
| OES Linux + SP1 | ZENworks 7 and 7 w/SP1 only | 2 | 6 | <i>variable</i> ² |
| OES Linux + SP2 | ZENworks 7 SP1 only | 2 | 6 | <i>variable</i> ² |
| Red Hat Linux 7.1, 7.2, 7.3, 8 | ZfS 3.0.2 only | 2 | 4 | <i>variable</i> ² |
| Red Hat Linux 9 | ZENworks 6.5 only | 2 | 4 | <i>variable</i> ² |
| Red Hat Advanced Server 2.1 | ZENworks 7, ZENworks 7 w/SP1, ZENworks 6.5, ZENworks 6.5 SP1, ZENworks 6.5 SP2, or ZfS 3.0.2 | 2 | 4 | <i>variable</i> ² |
| Red Hat Enterprise Server 2.1 | ZENworks 7, ZENworks 7 w/SP1, ZENworks 6.5, ZENworks 6.5 SP1, ZENworks 6.5 SP2, or ZfS 3.0.2 | 2 | 4 | <i>variable</i> ² |
| Red Hat Enterprise Linux AS 3 | ZENworks 7, ZENworks 7 w/SP1, ZENworks 6.5, ZENworks 6.5 SP1, or ZENworks 6.5 SP2 | 2 | 4 | <i>variable</i> ² |
| Red Hat Enterprise Linux AS 4 | ZENworks 7 and 7 w/SP1 only | 2 | 4 | <i>variable</i> ² |
| Red Hat Enterprise Linux ES 3 | ZENworks 7, ZENworks 7 w/SP1, ZENworks 6.5, ZENworks 6.5 SP1, or ZENworks 6.5 SP2 | 2 | 4 | <i>variable</i> ² |
| Red Hat Enterprise Linux ES 4 | ZENworks 7 and 7 w/SP1 only | 2 | 4 | <i>variable</i> ² |
| Solaris 8 | ZfS 3.0.2 only | 5 | 8 | N/A |
| Solaris 9 | ZENworks 7, ZENworks 7 w/SP1, ZENworks 6.5, ZENworks 6.5 SP1, or ZENworks 6.5 SP2 | 5 | 9 | N/A |
| SUSE [®] Linux 8.1 ³ | ZfS 3.0.2 only | 2 | 4 | <i>variable</i> ² |
| SUSE Linux 8.2 | ZfS 3.0.2 only | 2 | 4 | <i>variable</i> ² |
| SUSE Linux Enterprise Server (SLES) 8 | ZENworks 7, ZENworks 7 w/SP1, ZENworks 6.5, ZENworks 6.5 SP1, or ZENworks 6.5 SP2 | 2 | 4 | <i>variable</i> ² |
| SUSE Linux Standard Server (SLSS) 8 | ZENworks 7, ZENworks 7 w/SP1, ZENworks 6.5, ZENworks 6.5 SP1, or ZENworks 6.5 SP2 | 2 | 4 | <i>variable</i> ² |
| SLES 9 | ZENworks 7, ZENworks 7 w/SP1, ZENworks 6.5, ZENworks 6.5 SP1, or ZENworks 6.5 SP2 | 2 | 6 | <i>variable</i> ² |
| SLES 9 SP1, SP2 | ZENworks 7 and 7 w/SP1 only | 2 | 6 | <i>variable</i> ² |
| SLSS 9 | ZENworks 7, ZENworks 7 w/SP1, ZENworks 6.5, ZENworks 6.5 SP1, or ZENworks 6.5 SP2 | 2 | 6 | <i>variable</i> ² |

| Operating System Version | Subscriber Version ¹ | Major | Minor | Revision |
|--------------------------|---------------------------------|-------|-------|------------------------------|
| SLSS 9 SP1, SP2 | ZENworks 7 and 7 w/SP1 only | 2 | 6 | <i>variable</i> ² |

¹ The Subscriber column indicates the Subscriber version that is required for processing the software package on a server platform. You do not need to specify the Subscriber version here; however, the software package cannot be successfully sent and extracted on a server with one of these network operating systems unless the correct product version for the Subscriber software is running on it.

It is possible to have both ZENworks 6.5 (or later) Server Management and ZENworks for Servers 3.0.2 running in your network, such as when you are upgrading incrementally. This table provides platform information for both ZENworks product versions. For information on upgrading incrementally, see “[Upgrade Concepts and Issues](#)” in the *Novell ZENworks 7 Server Management Installation Guide*.

² Where *variable* is listed for the Linux platforms in this table, the number that you enter in the *Revision* field is the kernel revision that was either first shipped with the installed Linux operating system, or a later version that you might have updated your Linux servers to. You can use the `uname -a` command to determine the exact revision number.

Also, the revision number you need to enter depends on what you specify in the *Version* field (*Equal To*, *Greater Than*, and so on). For example, you could enter 0 in the *Revision* field and select *Greater Than* in the *Version* field to include all kernel revisions.

³ Because Linux versions are the same for Red Hat and SUSE, and you can only select `Linux` in the *Platform* field, to differentiate between those two Linux distributions, elsewhere in the software package you can require a certain file belonging to either Red Hat or SUSE to exist on the server. For example, the `/etc/SuSE-release` file could be required on the server, so that only servers with the SUSE Linux version would be accepted for receiving the software package.

E.2 Memory (RAM)

To configure the server memory requirement:

- 1 With the memory requirement selected, select the condition:

- Less Than
- Less Than or Equal To
- Greater Than
- Greater Than or Equal To

- 2 Specify the size in megabytes of RAM for the condition selected.

E.3 Disk Space

To configure the disk space requirement:

- 1 With the disk space requirement selected, select the root location.

The two options are *SYS Volume* and *Volume*. To conserve disk space usage on NetWare servers, do not select the `sys :` volume if you have other volumes with available disk space.

IMPORTANT: Do not use literal volume/drive names, such as the forward slash (/) character, when you are sending to multiple platforms. For example if you specify the / character for Linux in a software package, and the Software Package Distribution also gets sent to Windows servers, Windows will recognize the / as meaning its root location and the files intended only for the Linux servers would be installed on the Windows servers. To solve this problem, do one of the following:

a. Use variables for disk space locations, which allows you to send to multiple platforms.

or

b. Use the Operating System requirement in conjunction with the Disk Space requirement to confine the Distribution to only the server platforms where the literal disk space location exists.

Examples of literal locations you can provide:

NetWare:

```
sys:  
data:
```

Windows:

```
C:\  
\\myserver\data\shared_folder
```

Linux or Solaris:

```
/  
/usr  
/usr/data  
/usr/data  
/etc  
/mnt/files
```

For Linux and Solaris servers, it is any path that identifies a disk partition.

2 If you selected *Volume*, provide the volume's name.

3 Select the condition:

- Less Than
- Less Than or Equal To
- Greater Than
- Greater Than or Equal To

4 Specify the free disk space needed in megabytes for the condition selected.

It is important that the free disk space you specify exists at the location you specified in [Step 1](#).

E.4 SET Commands

When adding SET commands, the SET Commands Wizard is automatically run.

To configure the SET commands requirement:

1 With the SET commands requirement selected, provide the name of the SET command.

2 Provide the SET command's value.

E.5 Registry

You can require that certain entries must exist in the registry before installing the software package.

To configure the registry requirement:

- 1 With the registry requirement selected, select the *Entry Type*:

- Key
- Name
- Data

- 2 For both entry types *Key* and *Name*, select if it exists or does not exist.

or

For the entry type *Data*, select if it equals or does not equal.

- 3 Enter the text for *Key*, *Name*, or *Data* (depending on which you selected in [Step 1](#)).

Make sure you add the two backslashes to the beginning of the *Key*. For example, `\\HKEY_LOCAL_MACHINE\software\....`

IMPORTANT: The % symbol is not valid in NetWare registry names.

HKEY_LOCAL_MACHINE does not convert to “My Server” for this *Registry* field entry as it does for the *Registry Settings* field (see [“Registry Settings” on page 265](#)).

E.6 File

To configure the file requirement:

- 1 With the file requirement selected, provide the name.

Include the file’s full path.

- 2 Select the required file status:

- File Exists
- File Does Not Exist
- Date Is

E.7 Products.dat

WARNING: Modifying the `products.dat` file could prevent something from running or being installed on the NetWare server. Never modify any entries supplied by Novell®.

To configure the `products.dat` requirement:

- 1 With the `products.dat` requirement selected, provide the name of item in the `.dat` file.

IMPORTANT: Names are case sensitive.

The item is the ID of the product in the `.dat` file.

- 2 Provide the version text that corresponds with the item selected in [Step 1](#).
- 3 Select whether the version *Contains*, *Begins With*, or *Matches* the version specified in [Step 2](#).

- 4 Provide the description text that corresponds with the item selected in **Step 1**.
- 5 Select whether the description *Contains*, *Begins With*, or *Matches* the description provided in **Step 4**.

Registry Entries for Server Software Package Components

F

The following information is used in several setup steps for software packages. For more information, see [“Registry Settings” on page 265](#).

The NetWare® or Windows registry entries you can change are keys, value names, and value data. You can select keys and value data types for making changes, and you can provide the corresponding value names when you select one of the types.

In all cases, you must enter the exact key name or value name that is expected in the registry, as well as the correct data values.

The registry settings under HKEY_LOCAL_MACHINE are the only ones you can configure using a software package.

You can change the following registry entries when you install a software package:

- ◆ [Section F.1, “Key,” on page 421](#)
- ◆ [Section F.2, “Binary,” on page 422](#)
- ◆ [Section F.3, “Expand String,” on page 422](#)
- ◆ [Section F.4, “\(Default\),” on page 422](#)
- ◆ [Section F.5, “DWord,” on page 423](#)
- ◆ [Section F.6, “Multi-Value String,” on page 423](#)
- ◆ [Section F.7, “String,” on page 423](#)

F.1 Key

Keys create the paths to the various registry entries. For example, HKEY_LOCAL_MACHINE is a registry key at the root level, and HARDWARE is a key directly under it. The keys are displayed with folder icons in tree fashion. You can click the plus or minus signs to expand or compress the tree structure.

In the box where the HKEY_LOCAL_MACHINE key is displayed, you need to use the Key registry entry to create the path to where the registry changes are placed.

To configure a Key entry:

- 1** In the box displaying your key tree, select the location where you want the key inserted.
- 2** Select *Key* from the drop-down box, then click *Add*.
New Key #1 is displayed.
- 3** Change the default key name to the key name that you need.
When entering information into this field, you must press Enter for the change to be saved.
- 4** Select a condition for making the registry change:
 - Create
 - Delete

- 5 To apply the setting to all subordinate keys, click *Apply To All*.

F.2 Binary

A value data type that is a list of hexadecimal numbers, such as:

```
d0 04 72 6e
```

You must first use the Key registry setting option to create the path to the key that holds the Binary information.

To configure a Binary entry:

- 1 In the box displaying your key tree, select the location where you want the binary data inserted.
- 2 Select *Binary* from the drop-down box, then click *Add*.

New Value #1 is displayed.

- 3 Change the default binary name to the name that you need.
- 4 Select a condition for making the registry change:

- Create
- Delete

- 5 Provide the binary data.

The *Data* box is a hexadecimal editor. There are three unlabeled columns:

First: Binary counter of the number of hexadecimal characters, beginning with 0000.

Second: Hexadecimal data, eight entries per row.

Third: Plain text ASCII characters corresponding to the hexadecimal data.

You can enter data in either the second or third column. As you enter data in one the second (hexadecimal) column, the corresponding characters are displayed in the third (text) column, and vice versa.

F.3 Expand String

NetWare only. Currently not supported.

F.4 (Default)

This is usually the first data entry for a key.

You must first use the Key registry setting option to create the path to the key that holds the (Default) entry.

To configure a (Default) entry:

- 1 In the box displaying your key tree, select the location where you want the (*Default*) entry made.
- 2 Select (*Default*) from the drop-down box, then click *Add*.
(Default) is displayed.
- 3 With the (*Default*) entry selected, select a condition for making the registry change:

Create
Delete

- 4 Enter a string in *Data*.

F.5 DWord

DWords are based on hexadecimal code that is represented in Double WORD format. For example:
0x00100022

You must first use the Key registry setting option to create the path to the key that holds the DWord information.

To configure a DWord entry:

- 1 In the box displaying your key tree, select the location where you want the DWord entry made.
- 2 Select *DWord* from the drop-down box, then click *Add*.
New Value #1 is displayed.
- 3 Change the default DWord name to the name that you need.
- 4 Select a condition for making the registry change:

Create
Delete

- 5 Enter the DWord string in *Data*.

F.6 Multi-Value String

NetWare only. Currently not supported.

F.7 String

String values are easy-to-read sequences of words or numbers within quote marks.

You must first use the Key registry setting option to create the path to the key that holds the String information.

To configure a String entry:

- 1 In the box displaying your key tree, select the location where you want the string data inserted.
- 2 Select *String* from the drop-down box, then click *Add*.
New Value #1 is displayed.
- 3 Change the default string name to the name that you need.
- 4 Select a condition for making the registry change:

Create
Delete

- 5 Enter the string in *Data*.

Client Access in Linux



Tiered Electronic Distribution sends application objects and content to Linux servers using file systems such as Reiser, Ext2, Ext3, and NSS (OES Linux only). However, NCP share names used in a Desktop Application object are not directly supported on Linux servers because Tiered Electronic Distribution cannot read NCP share names.

Do one of the following to ensure client access to files on Linux servers:

- ♦ [Section G.1, “Using Samba,” on page 425](#)
Use this option if you are running Samba.
- ♦ [Section G.2, “Using NCP Shares,” on page 425](#)
Use this option if you are not running Samba.

G.1 Using Samba

If you are accessing a Linux Subscriber server (including OES Linux) using a Windows client, access can be provided via a Samba share. Desktop Application objects hosted on that server will use the Samba share name in the UNC patch of the Application object.

To use Samba for share recognition:

- 1** On your Linux server, install Samba.
Refer to your Linux documentation for instructions on installing and configuring Samba.
- 2** Configure Samba by including a Samba share name (such as SYS) with the following path:
`/usr/novell/share`
where *share* is the Samba share name, such as *sys*.
- 3** Make sure Samba is running on the server when you send Desktop Application Distributions.
When a Desktop Application Distribution is configured with a golden application containing a path such as `\\192.68.1.203\sys\firefox`, the *sys* volume is interpreted as a share on the Linux server, allowing the Firefox application to be written to the `/usr/novell/sys/firefox` directory.
- 4** Repeat [Step 1](#) through [Step 3](#) for each Linux Subscriber and Distributor server where access to a Samba share name is required.
Include servers that have a location where:
 - ♦ The application files are to be written to (Linux Subscribers)
 - ♦ The application files are to be distributed from (Linux Distributors)

G.2 Using NCP Shares

If you do not want Samba running on your OES Linux Subscriber server, you can ensure that a Desktop Application Distribution will succeed by including the NCP share names in the Samba

configuration file. Samba does not need to be running in order for Tiered Electronic Distribution to use the `smb.conf` file.

To edit the Samba configuration file:

- 1 On your OES Linux server, open the following configuration file in a text editor:

```
/etc/samba/smb.conf
```

- 2 Copy an existing share section to the end of the file and paste it as many times as you have NCP shares to add, then edit the copied sections to be:

```
[share1]
comment = for access to the NCP share
path = /usr/novell/share1
write list =@ntadmin root
force group = netadmin
create mask = 0644
directory mask = 0755

[share2]
comment = for access to the NCP share
path = /usr/novell/share2
write list =@ntadmin root
force group = netadmin
create mask = 0644
directory mask = 0755
```

where `share1` and `share1` are the NCP share names, such as `sys` and `apps`. Note that the share names are used both within the brackets and in the path. The share names in the `smb.conf` file and in NCP must be the same.

NOTE: The last four lines for each new share section do not need to be edited from whatever is contained in the original copied section. The values shown above in these lines are not used by Tiered Electronic Distribution.

To display which NCP shares are available, enter `ncpcon` on the OES Linux server console and type `help` or `--h` for a list of the NCP command line options.

- 3 Save the file and exit the text editor.

When a Desktop Application Distribution is configured with a golden application containing a path such as `\\192.68.1.203\sys\firefox`, the `sys` volume is interpreted as a share on the OES Linux server, allowing the Firefox application to be written to the `/usr/novell/sys/firefox` directory.

- 4 Identify the servers that have a location where:

- ♦ The application files are to be written to (Subscribers)
- ♦ The application files are to be distributed from (Distributors)

then do one of the following:

- ♦ If you have only a few servers to configure, repeat **Step 1** through **Step 3** for each Linux Subscriber and Distributor server where access to an NCP share name is required.
- ♦ If you have very many servers that need this configuration change, use a Text File Changes policy (in the Distributed Server Package) to roll out an updated `smb.conf` file to them. For more information, see **“Text File Changes” on page 229**.

Configuration Planning Worksheet



Use the following worksheet to log configuration information as you plan how to set up your distribution system. You might need to attach lists for some items.

This worksheet is designed to print best from the PDF version of the documentation.

IMPORTANT: Do not use this planning worksheet by itself to configure Policy and Distribution Services, even if you feel experienced enough to do so. There are some required configuration steps that are not covered in this worksheet, because planning is not needed for those steps. Use the sections under [Section 1.2, “Configuring Your Distribution System,” on page 50](#) as your guide for performing the actual configuration of Policy and Distribution Services.

| Configuration Information | Instructions |
|---------------------------|--------------|
|---------------------------|--------------|

| | |
|---|--|
| Installing Additional Distributors, Databases, and Subscribers | |
|---|--|

| | |
|--|---|
| | If you do not have additional Distributors, databases, or Subscribers to install, skip to worksheet item 12 . |
|--|---|

| | |
|--|--|
| 1) Tree for the Distributor and ZENworks Database objects: | |
|--|--|

| | |
|--|---|
| | Provide the name of the eDirectory tree for installing the Server Management objects. |
|--|---|

For more information, see [Section 1.1.3, “Understanding Your Network Topology,” on page 36](#).

| | |
|------------------------------|--|
| 2) Distributor server names: | |
|------------------------------|--|

| | |
|--|--|
| | Provide the server names for each server that you want to be a Distributor. Distributor servers build and own the Distributions. |
|--|--|

For more information, see [“Distributor Properties” on page 38](#).

| | |
|-----------------------------|--|
| 3) Subscriber server names: | |
|-----------------------------|--|

| | |
|--|--|
| | Provide the server names for each server that you want to be a Subscriber. Subscriber servers receive and extract the Distributions. |
|--|--|

For more information, see [Section 1.1.5, “Other Subscribers To Be Installed?,” on page 41](#).

| Configuration Information | Instructions |
|--|--|
| 4) Database server names: | <p>Provide the server names for each server where you want to install the Server Management database, which can be installed on NetWare and Windows servers.</p> <p>You can have multiple databases for Policy and Distribution Services, but only one per server.</p> <p>Also specify the purpose for each database, or a Distributor identifier for each database if they will each be used the same way.</p> <p>For more information, see “Whether a Distributor Server Will Host a Server Management Database” on page 39.</p> |
| 5) Installation paths for Distributors’ software: | <p>Provide the path where you want the Distributor software installed. The default is <code>\zenworks</code> for both NetWare and Windows servers.</p> <p>For more information, see “Software Installation Paths” on page 39.</p> |
| 6) Installation paths for Subscriber software: | <p>Provide the path where you want the Subscriber software installed. The default is <code>\zenworks</code> for both NetWare and Windows servers.</p> <p>For more information, see “Software Installation Paths” on page 39.</p> |
| 7) Distributors’ properties, where different than the installation defaults: | <p>Edit the following information for your Distributor servers:</p> <ul style="list-style-type: none"> ◆ Distributor object’s name (the default is <code>Distributor_server_name</code>) ◆ Distributor’s context ◆ Distributor server’s working directory <p>For more information, see “Distributor Properties” on page 38.</p> |

| Configuration Information | Instructions |
|--|--|
| <p>8) Subscribers' properties, where different than the installation defaults:</p> | <p>Edit the following information for your Subscriber servers:</p> <ul style="list-style-type: none"> ◆ Subscriber object's name (the default is <code>Subscriber_server_name</code>) ◆ Subscriber context ◆ Subscriber server's working directory <p>For more information, see Section 1.1.5, "Other Subscribers To Be Installed?", on page 41.</p> |
| <p>9) Installation paths for Server Management database software:</p> | <p>Provide the path where you want the <code>zfslog.db</code> file located. The default is <code>\zenworks\database</code>. For NetWare servers, we recommend not using the <code>sys:</code> volume because the database file can become very large. We also recommend that you install the database software on a server where the Subscriber software is also installed so that you can use the Database Purge option.</p> <p>For more information, see "Software Installation Paths" on page 39.</p> |
| <p>10) Database object name:</p> | <p>Either accept the default names, or provide ones that will help you to identify the databases' purposes.</p> <p>For more information, see "Whether a Distributor Server Will Host a Server Management Database" on page 39.</p> |

| Configuration Information | Instructions |
|--|--|
| 11) Database object Container: | <p>We recommend you use the same container where your other Tiered Electronic Distribution objects reside.</p> <p>For more information, see “Whether a Distributor Server Will Host a Server Management Database” on page 39.</p> |
| Configuring the Distributors for a Mixed eDirectory Environment | <p>If you do not have a mixed eDirectory environment, skip to worksheet item 13.</p> |
| 12) IP address of server in eDirectory 8.x: | <p>Provide the IP address of a server in the tree using eDirectory 8.x. This can be the Distributor server’s IP address, if that server is running eDirectory 8.x.</p> <p>For more information, see “Whether Distributors Might Exist in a Mixed eDirectory Environment” on page 40.</p> |
| Installing Inter-Server Communications | <p>If you do not need to set up inter-server communications, skip to worksheet item 14.</p> |
| 13) Subscriber servers outside your secured network: | <p>Inter-server communications security might be needed if your Distributor and Subscriber servers communicate with servers outside your secured network.</p> <p>For more information, see “Determining Whether You Need Inter-Server Communications Security” on page 44.</p> |

| Configuration Information | Instructions |
|--|--|
| Installing NCI on Windows, Linux, or Solaris Servers | If you do not need to install NCI to these servers, skip to worksheet item 15 . |
| 14) Windows, Linux, or Solaris servers (Distributor or Subscriber) that will be involved with Distribution encryption: | <p>List the Windows, Linux, or Solaris servers that will either build (Distributors) or extract (Subscribers) encrypted Distributions.</p> <p>For more information, see “Determining Whether You Need Encryption Security for Windows Servers” on page 45.</p> |

Configuring the Distributor Routing Hierarchies

| | |
|--|--|
| 15) Distributors' routing hierarchies of tiered Subscribers: | <p>Create a chart of tiered Subscribers for each Distributor showing how you want your Distributions to be distributed on your network. Distributors can use Subscribers in other Distributor's routing hierarchies. However, a Subscriber should only be used once in a given Distributor's hierarchy so that an end-node Subscriber only has one distribution path for receiving a particular Distribution.</p> <p>For more information, see Section 1.1.6, “Determining the Distribution Flow,” on page 41.</p> |
|--|--|

| Configuration Information | Instructions |
|---------------------------|--------------|
|---------------------------|--------------|

Configuring Parent Subscribers

| | |
|--|--|
| 16) Subscriber/parent Subscriber assignments (end-node Subscribers associated with a parent Subscriber): | <p>Create Subscriber lists where each parent Subscriber delivers Distributions. You should assign each end-node Subscriber to a parent Subscriber, except where you want the end-node Subscriber to receive its Distribution directly from the Distributor.</p> <p>For more information, see "Selecting Subscribers for the Distribution Routes" on page 43.</p> |
|--|--|

Creating and Configuring Subscriber Groups

| | |
|--------------------------------------|---|
| 17) Subscriber Group object name: | <p>If you are not using Subscriber Groups, skip to worksheet item 19.</p> <p>Provide a unique name for the Subscriber Group.</p> <p>For more information, see "Subscriber Groups" on page 47.</p> |
| 18) Subscribers to be in this group: | <p>Provide a list of Subscribers that need the same Distributions from the Channel where the group is subscribed.</p> |

Configuration Information**Instructions**

Creating the Policy Package Distributions

19) Distributions, their types, and their Distributors:

Create a list of your Distributions. For each Distribution, include the Distribution type, object name, and servers that need the Distribution. The Distribution types are:

- ◆ “Desktop Application” on page 32
- ◆ “File” on page 33
- ◆ “FTP” on page 33
- ◆ “HTTP” on page 34
- ◆ “MSI” on page 34
- ◆ “Policy Package” on page 34
- ◆ “RPM” on page 36
- ◆ “Software Package” on page 36

For more information, see [Section 1.1.2, “Selecting Your Distributions,”](#) on page 32.

| Configuration Information | Instructions |
|--|--|
| Creating and Configuring the Channels | |
| 20) eDirectory container for Tiered Electronic Distribution objects: | <p>Container for creating and managing Tiered Electronic Distribution objects.</p> <p>You might have created a container during installation of Policy and Distribution Services. If not, you should create a container specifically for managing Tiered Electronic Distribution objects.</p> <p>For more information, see “Whether a Distributor Server Will Host a Server Management Database” on page 39.</p> |
| 21) Channel names: | <p>Provide the names of the Channel objects that you need for your Distributions. We recommend a unique Channel for each unique Distribution or Distribution grouping.</p> <p>For more information, see Section 1.1.8, “Determining the Channels for the Distributions,” on page 46.</p> |
| 22) Distributions for the Channels: | <p>Create a list of which Distributions belong to which Channels.</p> <p>For more information, see Section 1.1.8, “Determining the Channels for the Distributions,” on page 46.</p> |

Configuration Information**Instructions**

Subscribing to the Channels

23) Subscribers' Extract schedules:

Set extract schedules per Subscriber server according to when it would be best for each Subscriber to be extracting its Distributions.

For more information, see [Section 1.1.9, "Determining Subscribers' Subscriptions," on page 47.](#)

24) Channel associations with Subscribers and Subscriber Groups:

Create lists where Subscribers and Subscriber Groups are associated with the Channels that have the Distributions you want them to receive.

Documentation Updates

This section contains information on documentation content changes that were made in this *Administration Guide* after the initial release of Novell® ZENworks® 7 Server Management. The information can help you to keep current on updates to the documentation.

All changes that are noted in this section are also made in the documentation. The documentation is provided on the Web in two formats: HTML and PDF. The HTML and PDF documentation are both kept up-to-date with the documentation changes listed in this section.

The documentation update information is grouped according to the date the changes are published. Within a dated section, the changes are alphabetically listed by the names of the main table of contents sections in the guide.

If you need to know whether a copy of the PDF documentation you are using is the most recent, the PDF document contains its publish date on the front title page.

The documentation was updated on the following dates:

- ◆ [Section I.1, “March 29, 2007,” on page 437](#)
- ◆ [Section I.2, “August 16, 2006,” on page 438](#)
- ◆ [Section I.3, “July 14, 2006 \(Support Pack 1\),” on page 439](#)
- ◆ [Section I.4, “December 9, 2005,” on page 441](#)
- ◆ [Section I.5, “October 7, 2005,” on page 442](#)

I.1 March 29, 2007

Updates were made to the following sections:

| Location | Change |
|---|---|
| “Determining the Distributor’s Refresh Schedule” on page 49 | <p>Updated this section to clarify use of the Refresh schedule. For example, the default of Never is recommended because an infinite loop situation could be caused by a Refresh schedule that occurs more frequently than the time it takes to build or send a Distribution.</p> <p>Also updated the following sections with information related to this change:</p> <ul style="list-style-type: none">◆ “Scheduling the Distribution and Refreshing the Distributor” on page 60◆ Section 3.2.2, “The Basic Distribution Process,” on page 88◆ “Scheduling” on page 96◆ Section 3.3.5, “Manually Refreshing the Distributor,” on page 109 |
| “File” on page 118 | <p>Added the following paragraph to the IMPORTANT note in the section:</p> <p>Also, if a NetWare server is the target for a File Distribution, you might encounter an error due to code page differences where extended characters are used (such as ê, ë, ì, or í). The information in “Extended Characters in Directory Paths” on page 287 in the Desktop Application Distribution section is also applicable to File Distributions.</p> |

| Location | Change |
|---|---|
| "MSI" on page 120 | Added the following note to the section: IMPORTANT: Because an MSI Distribution recursively gathers all of the files from the MSI file's location, if you have multiple .msi files in a given location, all other files and subdirectories contained therein are gathered once for each .msi file. The distribution gathering process cannot determine which other files or subdirectories belong to each .msi file, so you can end up with a much larger MSI Distribution file than is necessary. Therefore, instead of storing your .msi files in one directory, place each into its own subdirectory with its own supporting files and subdirectories. |
| Section 3.10.5, "Sending Distributions: Firewall and Cluster Issues," on page 175 | Added this section concerning distribution issues for firewalls and clusters. |

I.2 August 16, 2006

Updates were made to the following sections:

- ♦ [Load/Unload Actions](#)
- ♦ [Security in Policy and Distribution Services](#)
- ♦ [Server Policies](#)
- ♦ [Tiered Electronic Distribution](#)

I.2.1 Load/Unload Actions

The following changes were made in this section:

| Location | Change |
|---|---|
| Section D.2, "Load Java Class," on page 411 | Updated this section with example information to make it more clear how to use the Load Java Class feature. |

I.2.2 Security in Policy and Distribution Services

The following changes were made in this section:

| Location | Change |
|--|---|
| Section 7.3.5, "TCP/IP Addresses and DNS Names," on page 315 | Removed the note concerning not using underscores in server names, as this restriction no longer applies. |

I.2.3 Server Policies

The following changes were made in this section: [Section 7.3.5, “TCP/IP Addresses and DNS Names,” on page 315](#)

| Location | Change |
|--|--|
| “Verifying Community String Changes” on page 225 | Added this new section which provides instructions on using TCPCON to view community string changes. |

I.2.4 Tiered Electronic Distribution

The following changes were made in this section:

| Location | Change |
|---|---|
| “Message Notification Levels” on page 183 | Updated this section, further explaining the messaging level types. |
| “Managing Message Notification Level Log Files” on page 184 | Created this section from information previously contained in “Message Notification Levels” on page 183 . |

I.3 July 14, 2006 (Support Pack 1)

Updates were made to the following sections:

- ◆ [Desktop Application Distribution](#)
- ◆ [Distribution Types](#)
- ◆ [Novell iManager](#)
- ◆ [Requirements for Server Software Packages](#)
- ◆ [Schedule Types](#)
- ◆ [Server Policies](#)
- ◆ [Tiered Electronic Distribution](#)

I.3.1 Client Access in Linux

The following changes were made in this section:

| Location | Change |
|---|---|
| Appendix G, “Client Access in Linux,” on page 425 | Added this new section which provides instructions on setting up client access on Linux servers (including OES Linux) to NCP share names used in Desktop Application objects. |

I.3.2 Desktop Application Distribution

The following changes were made in this section:

| Location | Change |
|---|---|
| Section 6.3.2, "Creating the Distribution," on page 290 | Removed documentation for the <i>Delete previous revision before receiving next</i> field, which is no longer used, and added documentation for the e-mail fields that replaced it. |

I.3.3 Distribution Types

The following changes were made in this section:

| Location | Change |
|---|--|
| Section A.2, "File," on page 383 | Updated the information in this section. |
| Section A.2.8, "Maintain Trustees," on page 387 | Added the following sentence: If synchronization is enabled for directories in a Distribution, the trustees of those directories are also synchronized. |
| Section A.3, "FTP," on page 387 | Updated the information in this section. |
| Section A.4, "HTTP," on page 390 | Updated the information in this section. |

I.3.4 Novell iManager

The following changes were made in this section:

| Location | Change |
|--|---|
| Section 2.4.1, "Setting Up Passwords for Remote Web Console," on page 73 | Added this section, which documents the addition of passwords for Remote Web Console in iManager. |
| Section 2.4.2, "Managing the Distributor Agent," on page 77 and Section 2.4.3, "Managing the Policy/Package Agent," on page 80 | Updated these sections with password entry steps. |

I.3.5 Requirements for Server Software Packages

The following changes were made in this section:

| Location | Change |
|--|---|
| Section E.1, "Operating System," on page 413 | Updated the table for the <i>Major</i> , <i>Minor</i> , and <i>Revision</i> fields with newer information and the OES platform. |

I.3.6 Schedule Types

The following changes were made in this section:

| Location | Change |
|---|--|
| Appendix B, "Schedule Types," on page 399 and Section B.5, "Never," on page 401 | Updated the information for the <i>Never</i> option. |

I.3.7 Server Policies

The following changes were made in this section:

| Location | Change |
|-----------------------------------|---|
| "Server Down Process" on page 221 | <p>Added the following note:</p> <hr/> <p>IMPORTANT: For the Windows, Linux, and Solaris platforms, if you down the server from its console, this policy is not recognized. Instead, you must down the server using the <i>Actions</i> option in <i>Remote Web Console</i> in iManager so that this policy can be applied.</p> |

I.3.8 Tiered Electronic Distribution

The following changes were made in this section:

| Location | Change |
|---|---|
| "Maximum Revisions" on page 115 and Section 3.4.4, "Creating a Distribution," on page 123 | Removed documentation for the <i>Delete previous revision before receiving next</i> field, which is no longer used, and added documentation for the e-mail fields that replaced it. |

I.4 December 9, 2005

Page design is reformatted to comply with revised Novell documentation standards.

I.5 October 7, 2005

Updates were made to the following sections:

- ♦ [Desktop Application Distribution](#)

I.5.1 Desktop Application Distribution

The following changes were made in this section:

| Location | Change |
|---|--|
| Section 6.3.2, "Creating the Distribution," on page 290 | In Step 7b , added information for the new option: Overwrite Existing Target Folder Object Attributes. |

Server Inventory



The Server Inventory component of Novell® ZENworks® 7 Server Management enables you to collect hardware and software inventory information from local and remote servers of your enterprise. This inventory information is scanned and stored in a database that can be accessed by the ZENworks administrator.

From Novell ConsoleOne®, you can view the complete hardware and software inventory of the servers. You can also query the centralized database of the servers.

The following sections provide information on the features and tasks of Server Inventory:

- ◆ [Chapter 12, “Understanding Server Inventory,” on page 445](#)
- ◆ [Chapter 13, “Setting Up Server Inventory,” on page 459](#)
- ◆ [Chapter 14, “Understanding the Server Inventory Components,” on page 529](#)
- ◆ [Chapter 15, “Understanding the ZENworks 7 Server Managements Inventory Database Schema,” on page 549](#)
- ◆ [Chapter 16, “Managing Your Inventory Information,” on page 581](#)
- ◆ [Chapter 17, “Viewing Inventory Information,” on page 635](#)
- ◆ [Chapter 18, “Monitoring Server Inventory Using Status Logs,” on page 711](#)
- ◆ [Appendix J, “Performance Tips,” on page 717](#)
- ◆ [Appendix K, “Hardware Information Collected by the Inventory Scanners,” on page 729](#)
- ◆ [Appendix L, “ZENworks 7 Server Management Inventory Attributes,” on page 747](#)
- ◆ [Appendix M, “Enumeration Values,” on page 767](#)
- ◆ [Appendix N, “Setting up Security for Server Inventory,” on page 777](#)
- ◆ [Appendix O, “Documentation Updates,” on page 779](#)

Understanding Server Inventory

12

The Server Inventory component of Novell® ZENworks® 7 Server Management gathers hardware and software inventory information from Novell NetWare® and Windows* servers in your enterprise and stores into a centralized database. Using this database, the network administrator can view and query for complete inventory information for the enterprise.

The inventory information can be useful to help you make business decisions on how to manage servers. The following are some of the business decisions that you can make once you have obtained the inventory information:

- ◆ Servers that need new applications
- ◆ Servers that need updated hardware and drivers
- ◆ Servers that conform to the corporate hardware and software standards

This chapter provides a basic overview of Server Inventory. It contains the following information:

- ◆ [Section 12.1, “Server Inventory Terminology,” on page 445](#)
- ◆ [Section 12.2, “Overview of Server Inventory Components,” on page 446](#)
- ◆ [Section 12.3, “Understanding Inventory Scanning Cycle,” on page 448](#)
- ◆ [Section 12.4, “Understanding the Inventory Server Roles,” on page 448](#)

12.1 Server Inventory Terminology

The following brief glossary provides basic definitions of Server Inventory terms:

Inventoried server: A server whose hardware and software information you want to scan and maintain in a central repository. To gather complete hardware and software inventory for a server, you must install the Inventory Agent on that server.

Inventory server: A server where you run the Inventory service. This server can run any other ZENworks 7 Server Management services also. The Inventory server collects the inventory information from a group of associated inventoried servers and stores it into the Inventory database. If you want to collect the inventory for the Inventory server, you must install the Inventory Agent on that Inventory server.

Inventory database: A repository of inventory information of all the inventoried servers.

Database server: A server running Sybase*, Oracle*, or MS SQL where your Inventory database is mounted. The database can run on an Inventory server or on a different server.

Management console: A Windows workstation or server running Novell ConsoleOne® with Server Inventory ConsoleOne snap-ins installed. The management console provides the interface to administer the inventory system.

eDirectory Tree: The Novell eDirectory™ tree consists of eDirectory objects such as multiple levels of organizational units, users, groups, and other network resources. This hierarchical structure is referred to as the eDirectory tree in this document. For more information, see the [Novell eDirectory documentation Web site \(http://www.novell.com/documentation\)](http://www.novell.com/documentation).

Inventory tree: A logical tree depicting the transmission of the inventory information from the inventoried servers and the Inventory servers to the centralized enterprise Inventory database.

Standalone Server: An Inventory server that has an Inventory database and inventoried servers attached to it.

Leaf Server: The lowest-level Inventory server in the inventory tree hierarchy. This server has one or more inventoried servers attached to it and can have an Inventory database attached to it. This Inventory server collects the inventory information from the inventoried servers attached to it and moves the information to the next-level Inventory server.

Intermediate Server: The Inventory server for moving the inventory information from the lower-level Inventory servers up the Inventory server hierarchy. This server can have either inventoried servers or an Inventory database, or both attached to it.

Root Server: The highest-level Inventory server in the inventory tree hierarchy. This server has the Inventory database that contains the inventory information of all the lower-level Inventory servers. At the Root Server level, you can view complete inventory information for the entire enterprise. This server can have inventoried servers attached to it.

Site: A site is typically a geographical location. There can be multiple sites your enterprise.

Software Dictionary or Dictionary: The software dictionary contains a list of software identifiers and rules. Each software identifier identifies a particular product installed on an inventoried server.

Software Identifiers: An entry that identifies a software product is called as a software identifier. Each software identifier has a set of file matching attributes and corresponding software information attributes. During the Inventory scan, the scanner reads the attributes from the file headers, and if these attributes match the attributes configured in the dictionary, the information in the corresponding software information attributes is stored in the Inventory database.

Software Dictionary Rule: A software dictionary rule represents a set of conditions that control the scope of scanning process.

Unidentified Software: The software dictionary might not identify all the software installed in your network. The software that is not listed in the dictionary is called Unidentified software.

12.2 Overview of Server Inventory Components

Before setting up the ZENworks 7 Server Inventory, you should understand the inventory components that interact together to perform inventory functions.

Server Inventory uses the following components:

- ◆ [Section 12.2.1, “Inventory Scanners,” on page 446](#)
- ◆ [Section 12.2.2, “Inventory Components on Inventory Servers,” on page 447](#)
- ◆ [Section 12.2.3, “Inventory Database,” on page 448](#)
- ◆ [Section 12.2.4, “Management Console,” on page 448](#)

12.2.1 Inventory Scanners

Platform-dependent scanners determine the hardware and software configurations of the inventoried servers. These scanners are located at the inventoried servers. When executed on the inventoried

servers, the scanners collect the inventory information. The inventory information is subsequently transferred to the Inventory server and processed.

Using the Server Inventory policy, you can configure the scan settings for scheduling the scan on the inventoried servers and customize hardware scanning. From the Inventory Service object (Inventory Service_ *server_name*), you can specify the location of the inventory information, and also customize software scanning using the Software Dictionary snap-ins.

NOTE: If you have inventoried servers that do not have the Novell Client installed, send their scans to the OES Linux Inventory server, then ensure that the OES server name is the same as the DNS name.

Also, the ZENworks Dekstop Management Linux Inventory server can receive scans from ZENworks Server Management inventoried servers.

For more information about the Inventory scanner, see [Section 14.3, “Understanding the Inventory Scanner,” on page 532](#).

12.2.2 Inventory Components on Inventory Servers

The Inventory server components process the inventory information. The following components are Java* programs that work identically on NetWare and Windows Inventory servers:

- ◆ Scan Collector

The Scan Collector collects the inventory information from the Inventory agent and stores them in an appropriate directory at the Inventory server. The inventory information is transferred using the XML-RPC protocol.

- ◆ Selector

The Selector processes the inventory information and places the information in appropriate directories. For more information, see [Section 14.5, “Understanding the Selector,” on page 544](#).

- ◆ Sender and Receiver

The Sender on the Inventory server compresses the inventory information and then transfers it from the lower-level Inventory server to the Receiver on the higher-level Inventory servers. By using the Roll-Up policy, you can configure the next level destination Inventory server for roll-up, and also schedule the roll-up time. For more information, see [Section 14.4, “Understanding the Sender and Receiver,” on page 539](#).

- ◆ Storer

The Storer stores the collected inventory information into the Inventory database. By using the Database Location policy, you can configure the properties of the Inventory Database object (Inventory database_ *server_name*) and associate the database object to an Inventory server. For more information, see [Section 14.6, “Understanding the Storer,” on page 545](#).

- ◆ Dictionary Provider and Dictionary Consumer

All Inventory servers run the Dictionary Provider and Dictionary Consumer services. The Dictionary Consumer downloads the dictionary updates from the Dictionary Provider. For more information, see [Section 14.7, “Understanding the Dictionary Provider and Dictionary Consumer,” on page 546](#).

12.2.3 Inventory Database

The Inventory database is a repository of inventory information of the inventoried servers. In Server Management, the database is a Common Information Model-based database and is implemented in Relational Database Management System (RDBMS). The database is maintained in Sybase, Oracle, or MS SQL. For more information, see [Section 13.2, “Setting Up the Inventory Database,” on page 493](#).

12.2.4 Management Console

The management console is the Novell ConsoleOne. This is a Java-based console that includes snap-ins for Server Inventory management operations.

12.3 Understanding Inventory Scanning Cycle

The Inventory scanning cycle is as follows:

1. The Inventory scanner checks whether an updated dictionary is available at its Inventory server and downloads the updated dictionary.
2. The Inventory scanner sends hardware and software information from the inventoried servers to the Inventory server as per the scan schedule.
3. The Inventory server stores the inventory information in the Inventory database.
4. At the management console, you can view and retrieve the inventory information from the Inventory database using Inventory tools such as Reporting, Summary, etc.

12.4 Understanding the Inventory Server Roles

This section describes the following roles that you can assign for an Inventory server:

- ◆ [Section 12.4.1, “Root Server,” on page 449](#)
- ◆ [Section 12.4.2, “Root Server with Inventoried Servers,” on page 450](#)
- ◆ [Section 12.4.3, “Intermediate Server,” on page 450](#)
- ◆ [Section 12.4.4, “Intermediate Server with Database,” on page 451](#)
- ◆ [Section 12.4.5, “Intermediate Server with Inventoried Servers,” on page 452](#)
- ◆ [Section 12.4.6, “Intermediate Server with Database and Inventoried Servers,” on page 453](#)
- ◆ [Section 12.4.7, “Leaf Server,” on page 454](#)
- ◆ [Section 12.4.8, “Leaf Server with Database,” on page 455](#)
- ◆ [Section 12.4.9, “Standalone Server,” on page 456](#)
- ◆ [Section 12.4.10, “Quick Reference Table of the Inventory Server Roles,” on page 457](#)

For a quick reference table of the Inventory Server roles, see [Section 12.4.10, “Quick Reference Table of the Inventory Server Roles,” on page 457](#).

12.4.1 Root Server

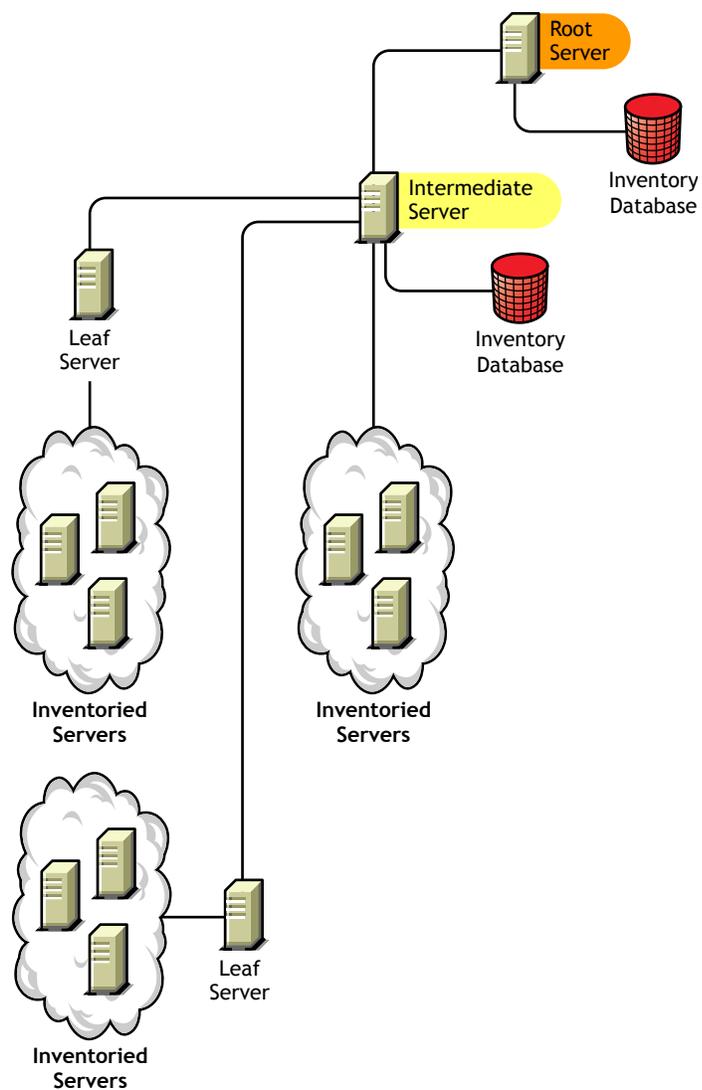
The Root Server has the following characteristics:

- ♦ This server is the topmost Inventory server in the inventory tree hierarchy.
- ♦ This server has an Inventory database attached to it.

Choose Root Server to store the inventory information for your enterprise in a centralized database. The Inventory database at the Root Server contains the inventory information for all the lower-level Inventory servers.

Figure 12-1 depicts Leaf Servers connected to the Intermediate Server with Database. The Intermediate Server is attached to the Root Server.

Figure 12-1 Root Server



12.4.2 Root Server with Inventoried Servers

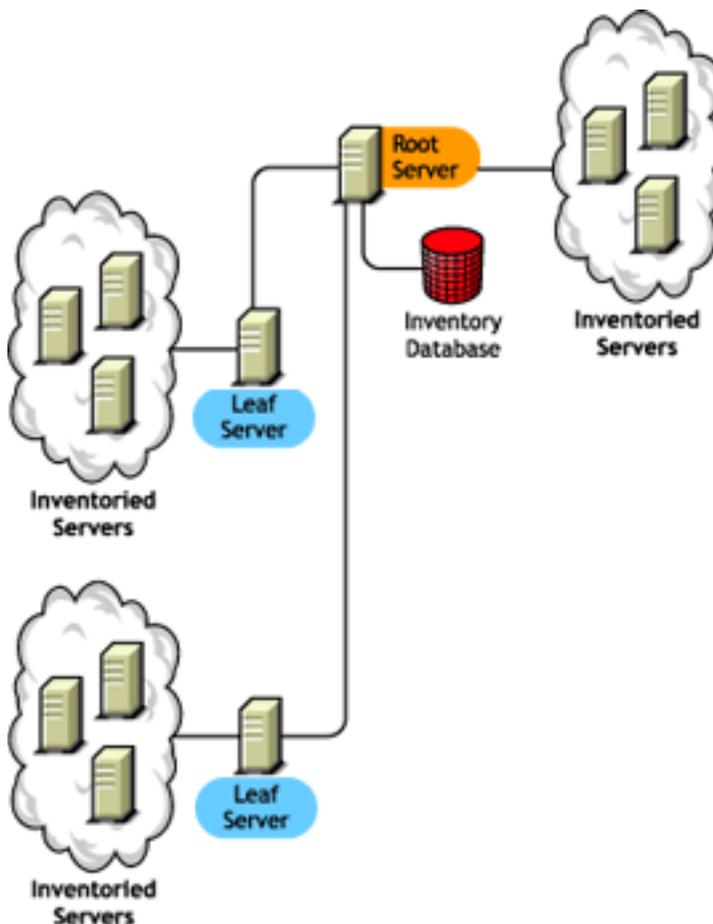
The Root Server with Inventoried Servers has the following characteristics:

- ♦ This server is the topmost Inventory server in the inventory tree hierarchy.
- ♦ This server has an Inventory database and inventoried servers attached to it. We recommend that you have to deploy these inventoried servers in a LAN.

Choose Root Server with Inventoried Servers if you want to store the inventory information of your enterprise in a centralized database and if you have inventoried servers in the same site as the Root Server. You can directly send the inventory information from these inventoried servers to the Root Server. The Inventory database at the Root Server with Inventoried Servers contains the inventory information for all these inventoried servers as well as for all the lower-level Inventory servers.

Figure 12-2 depicts a Root Server with inventoried servers and Inventory database attached to it. The Leaf Servers are connected to the Root Server.

Figure 12-2 Root Server with Inventoried Servers



12.4.3 Intermediate Server

The Intermediate Server has the following characteristics:

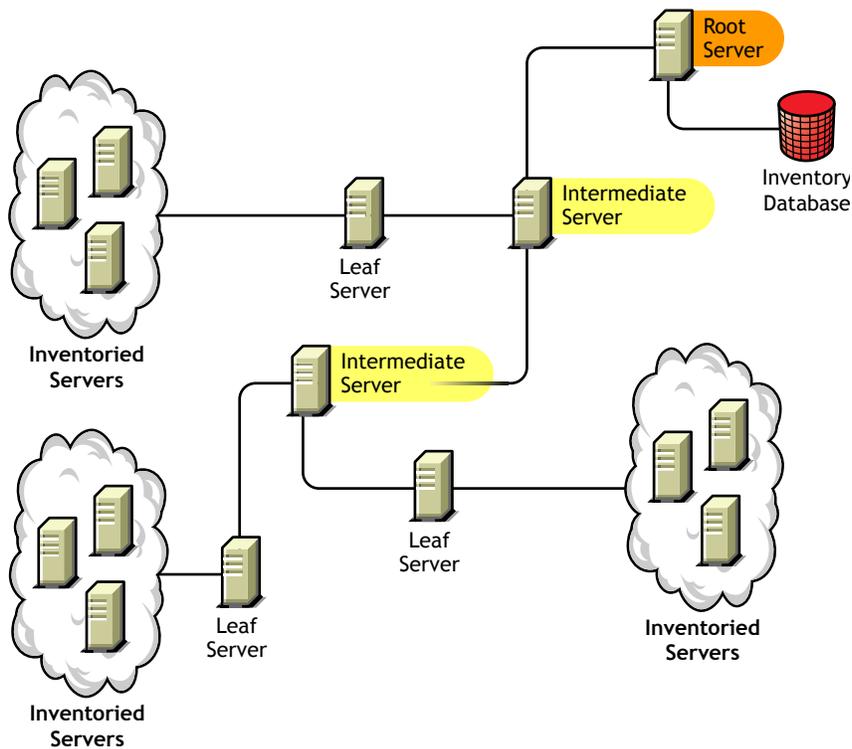
- ♦ This Inventory server acts as a staging server for the lower-level Leaf Servers.

- ◆ This server moves the inventory information to the next-level Inventory server.
- ◆ This server does not have inventoried servers or an Inventory database attached to it.
- ◆ There can be one or more Intermediate Servers in your enterprise.

Place Intermediate Servers on sites where the link parameters change substantially. The Intermediate Server stores the scan files to the disk to make up for the difference in bandwidth and reliability.

Figure 12-3 depicts an Intermediate Server connected to a Root Server. Two Leaf Servers roll up the inventory information to the Intermediate Server. This Intermediate Server rolls up the inventory information to another Intermediate Server that is connected to the Root Server.

Figure 12-3 Intermediate Server



In this illustration, there are many Leaf Servers and Intermediate Servers at different levels. The Intermediate Server is a staging server for uploading the scan information to the next-level server. The last Intermediate Server is attached to the topmost Root Server. This scenario is typical if there are many Leaf Servers in different geographical locations. All the Leaf Servers move the inventory information to the Intermediate Server.

In some scenarios, the Leaf Server connects to the Intermediate Server over a WAN.

12.4.4 Intermediate Server with Database

The Intermediate Server with Database has the following characteristics:

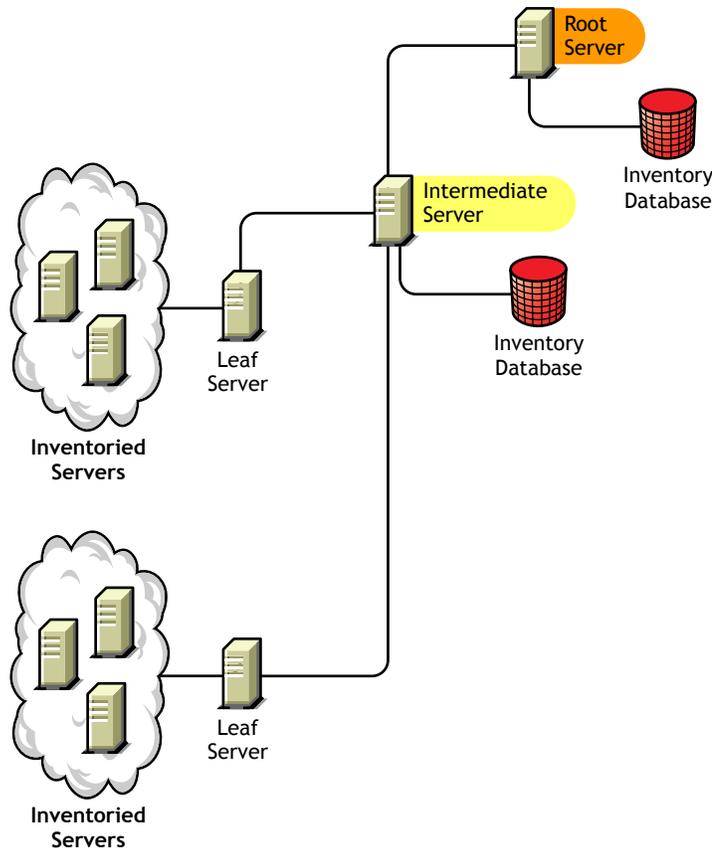
- ◆ This Inventory server acts as a staging server for the lower-level Leaf Servers.
- ◆ This server moves the inventory information to the next-level Inventory server.
- ◆ This server has an Inventory database attached to it.

- ◆ There can be one or more Intermediate Servers with Database in your enterprise.

Choose Intermediate Server with Database if you want to administer an intermediate site by generating Inventory reports. The inventory information that is rolled up to this Inventory server is stored in the local Inventory database and also rolled up to the next-level Inventory server.

Figure 12-4 depicts two Leaf Servers attached to the Intermediate Server. A consolidated inventory information of all Leaf Servers is available at the Intermediate Server level.

Figure 12-4 Intermediate Server with Database



12.4.5 Intermediate Server with Inventoried Servers

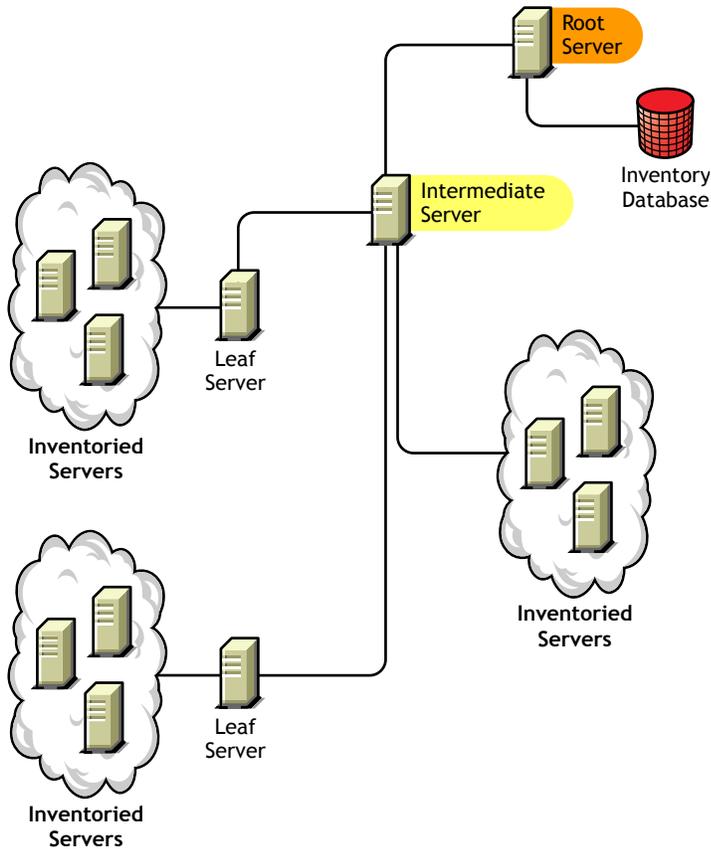
The Intermediate Server with Inventoried Servers has the following characteristics:

- ◆ This Inventory server acts as an intermediate server for the lower-level Leaf Servers.
- ◆ This server moves the inventory information to the next-level Inventory server.
- ◆ This server has inventoried servers attached to it.
- ◆ This server does not have an Inventory database attached to it.
- ◆ There can be one or more Intermediate Servers with Inventoried Servers in your enterprise.

Choose Intermediate Server with Inventoried Servers if you want an Intermediate Server and the site having the Intermediate Server has inventoried servers, whose inventory information you want to store it at the Root Server.

Figure 12-6 depicts two Leaf Servers attached to the Intermediate Server. This Intermediate Server also has inventoried servers attached to it.

Figure 12-5 Intermediate Server with Inventoried Servers



12.4.6 Intermediate Server with Database and Inventoried Servers

The Intermediate Server with Database and Inventoried Servers has the following characteristics:

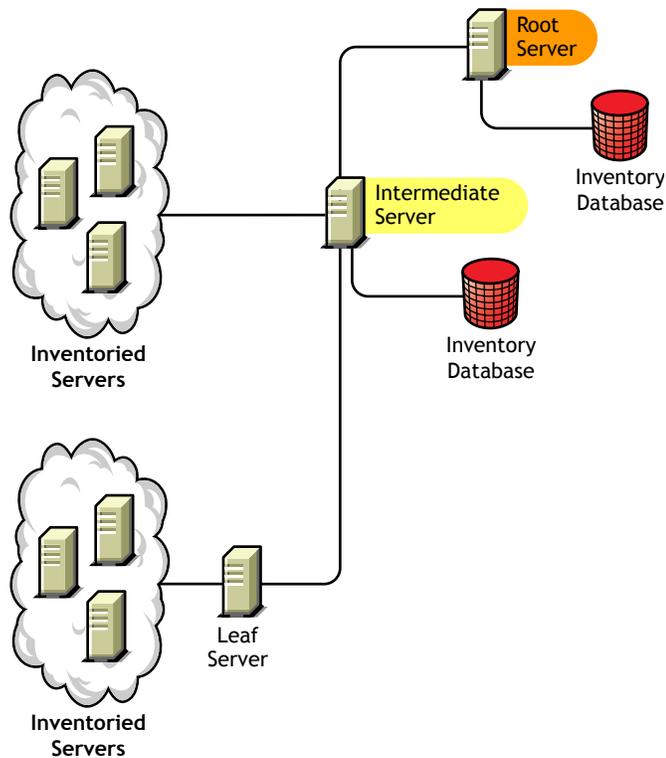
- ◆ This Inventory server acts as a staging server for the lower-level Leaf Servers.
- ◆ This server moves the inventory information to the next-level Inventory server.
- ◆ This server has inventoried servers attached to it.
- ◆ This server has an Inventory database attached to it.
- ◆ There can be one or more Intermediate Servers with Database and Inventoried Servers in your enterprise.

Choose Intermediate Server with Database and Inventoried Servers if you want the functionalities of [Intermediate Server with Database](#) and [Intermediate Server with Inventoried Servers](#) available on the site.

Figure 12-6 depicts two Leaf Servers attached to the Intermediate Server. The Intermediate Server has inventoried servers attached to it. A consolidated Inventory database of all Leaf Servers and the

inventoried servers that are directly connected to the Intermediate Server is available at the Intermediate Server level.

Figure 12-6 Intermediate Server with Database and Inventoried Servers



12.4.7 Leaf Server

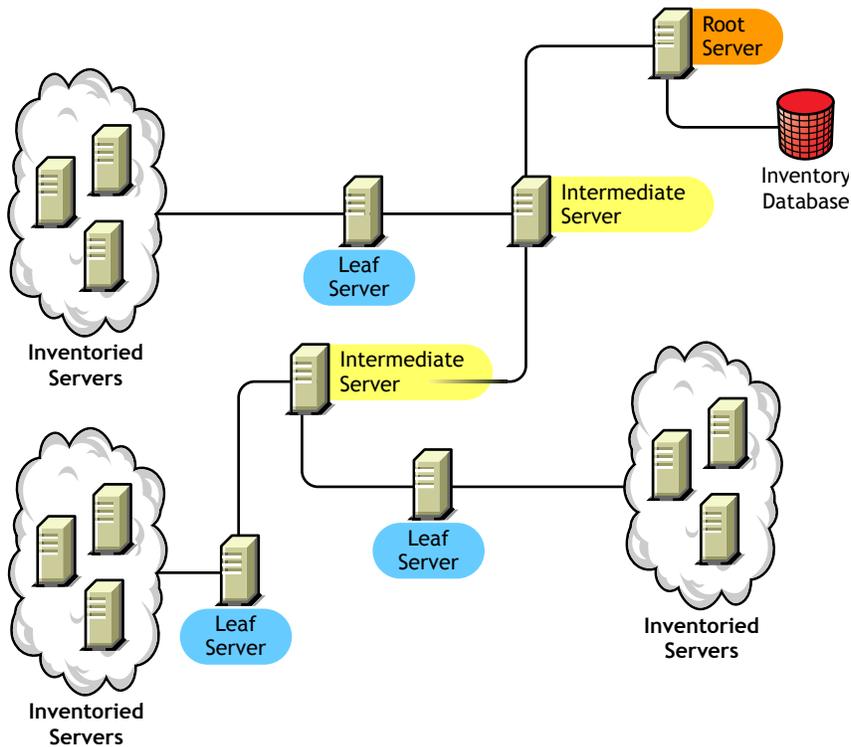
The Leaf Server has the following characteristics:

- ◆ This Inventory server is at the lowest level in the inventory tree hierarchy.
- ◆ This server has inventoried servers attached to it.
- ◆ This server moves the inventory information to the next-level Inventory server.
- ◆ A simple Leaf Server does not have an Inventory database. An Inventory database is not required because there might be only few inventoried servers attached to the Leaf server.

Choose Leaf Server if you have inventoried servers at remote sites, and you want to obtain and store the inventory information from these inventoried servers in a centralized database.

Figure 12-7 depicts many Leaf Servers attached to the Intermediate Server. The Intermediate Server is connected to a Root Server. A consolidated Inventory database of all Leaf Servers is available at the Root Server level.

Figure 12-7 Leaf Server



12.4.8 Leaf Server with Database

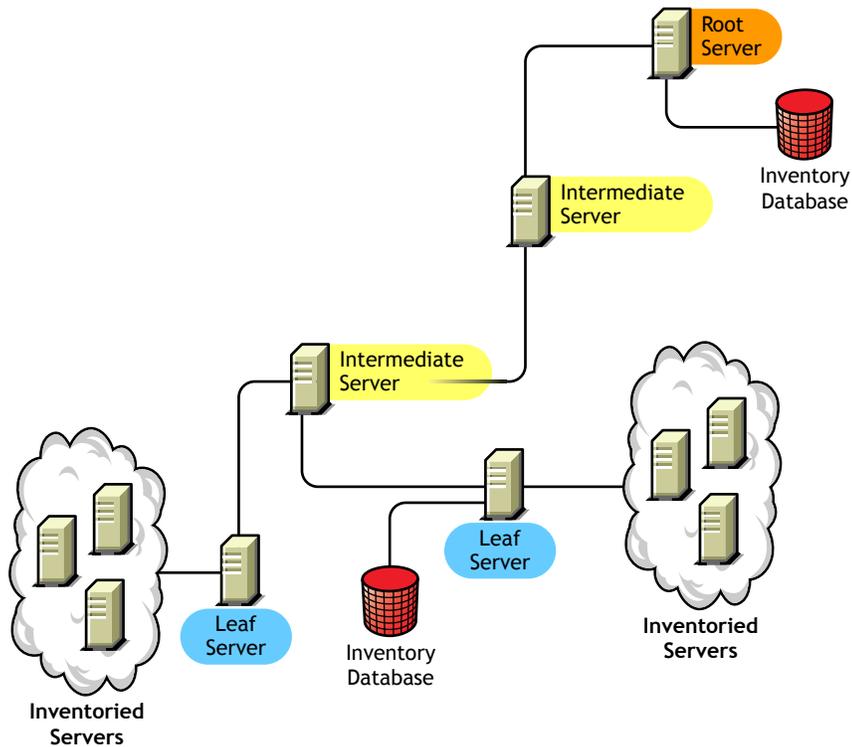
The Leaf Server with Database has the following characteristics:

- ◆ This Inventory server is at the lowest level in the inventory tree hierarchy
- ◆ This server has inventoried servers attached to it.
- ◆ This server moves the inventory information to the next-level Inventory server.
- ◆ This server has an Inventory database attached to it. You can assign a server as a Leaf Server with Database to maintain the inventory information for inventoried servers specific to the site.

Choose Leaf Server with Database if you want the functionalities of a **Leaf Server** as well as administer the site by generating Inventory reports

Figure 12-8 depicts two Leaf Servers attached to the Intermediate Server. One Leaf Server has an Inventory database attached to it. This database contains a consolidated inventory of all inventoried servers attached to this Leaf Server.

Figure 12-8 Leaf Server with Database



12.4.9 Standalone Server

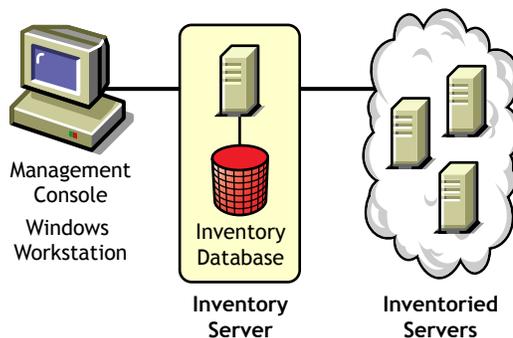
The Standalone Server has the following characteristics:

- ◆ This Inventory server has inventoried servers attached to it.
- ◆ This server has an Inventory database attached to it.
- ◆ There is no roll-up of scan information and there are no requirements for Intermediate Servers and a Root Server.

Use a Standalone Server if your network is made up of a single site and you want to administer that site.

Figure 12-9 depicts Standalone Server:

Figure 12-9 Standalone Server



12.4.10 Quick Reference Table of the Inventory Server Roles

Table 12-1 *Inventory Server Roles*

| Inventory Server | Is Inventory Database Attached to the Inventory Server? | Are Inventoried Servers Attached to the Inventory Server? |
|---|--|--|
| Root Server | Yes | No |
| Root Server with Inventoried Servers | Yes | Yes |
| Intermediate Server | No | No |
| Intermediate Server with Database | Yes | No |
| Intermediate Server with Inventoried Servers | No | Yes |
| Intermediate Server with Database and Inventoried Servers | Yes | Yes |
| Leaf Server | No | Yes |
| Leaf Server with Database | Yes | Yes |
| Standalone Server | Yes | Yes |

Setting Up Server Inventory

13

Before you install Novell® ZENworks® 7 Server Inventory in your working environment, you must plan and decide the hierarchy of the Inventory server tree for your enterprise. You should organize your inventory deployment based on your network constraints and information requirements.

The following sections contain detailed information to help you deploy Server Inventory in your enterprise:

- ◆ [Section 13.1, “Deploying Server Inventory,” on page 459](#)
- ◆ [Section 13.2, “Setting Up the Inventory Database,” on page 493](#)
- ◆ [Section 13.3, “Configuring the Inventory Service Object,” on page 520](#)
- ◆ [Section 13.4, “Configuring the Database Location Policy,” on page 521](#)
- ◆ [Section 13.5, “Configuring the Server Inventory Policy,” on page 522](#)
- ◆ [Section 13.6, “Configuring the Roll-Up Policy,” on page 524](#)
- ◆ [Section 13.7, “Configuring the Dictionary Update Policy,” on page 525](#)
- ◆ [Section 13.8, “Setting Up Distribution of Dictionary,” on page 526](#)

13.1 Deploying Server Inventory

The following sections help you to deploy Server Inventory:

- ◆ [Section 13.1.1, “Simple Deployment,” on page 459](#)
- ◆ [Section 13.1.2, “Advanced Deployment,” on page 463](#)
- ◆ [Section 13.1.3, “Understanding the Effects of Server Inventory Installation,” on page 480](#)
- ◆ [Section 13.1.4, “Starting and Stopping the Inventory Service,” on page 482](#)
- ◆ [Section 13.1.5, “Changing the Role of the Inventory Server,” on page 483](#)

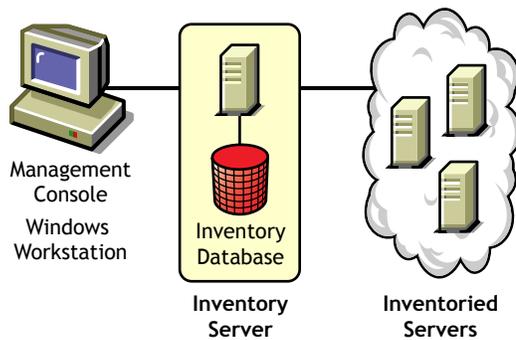
IMPORTANT: The recommendations discussed in the scenarios are generic. Based on the topology of your enterprise, further refinements might become necessary.

13.1.1 Simple Deployment

In the example scenario, the network consists of a single site and up to 5000 inventoried servers. The Inventory server components and the database are located on a Standalone Server, and the inventoried servers send scans to the Standalone server.

This scenario is illustrated in [Figure 13-1](#):

Figure 13-1 Simple Server Inventory Deployment



The following sections contain detailed information to help you deploy Server Inventory in a single site:

1. [“Recommendations for Deployment” on page 460](#)
2. [“Installing Server Inventory” on page 460](#)
3. [“Understanding the Effects of Server Inventory Installation” on page 461](#)
4. [“Configuring the Required Policies” on page 461](#)
5. [“Starting the Inventory Service” on page 461](#)
6. [“Updating the Software Dictionary” on page 461](#)
7. [“Understanding the Inventory Scanning Cycle in the Standalone Scenario” on page 461](#)

Recommendations for Deployment

- ◆ The minimum base Inventory server configuration includes 512 MB RAM and a database cache of 128 MB.
- ◆ The transmission of inventory information to the Inventory server and storage of the inventory information into the Inventory database is an ongoing backend process that can take several hours or even more than a day.
- ◆ If many inventoried servers are attached to the same Inventory server, we recommend that you do not schedule the scan of all inventoried servers at the same time, because this stresses Novell eDirectory™ and the Inventory services.
- ◆ Ensure that the eDirectory time synchronization radius is set within 2 seconds.
- ◆ The optimal database cache size requirement for the server could vary because of the server environment. Determine the database cache size that needs to be set by trying a range of cache sizes in the runtime environment. The default Sybase database cache size is 128 MB. For more information about improving the database performance, see [Appendix J, “Performance Tips,” on page 717](#).

Installing Server Inventory

During the Server Inventory installation, configure the Inventory Standalone Configuration settings. For detailed information on installing Server Inventory, see [“Policy-Enabled Server Management Installation”](#) in the *Novell ZENworks 7 Server Management Installation Guide*.

Understanding the Effects of Server Inventory Installation

For detail information on the effects of Server Inventory installation, see [Section 13.1.3, “Understanding the Effects of Server Inventory Installation,”](#) on page 480.

Configuring the Required Policies

Configure the [Server Inventory Policy](#).

Starting the Inventory Service

After installing ZENworks 7 Server Management, the Inventory service is automatically started.

Updating the Software Dictionary

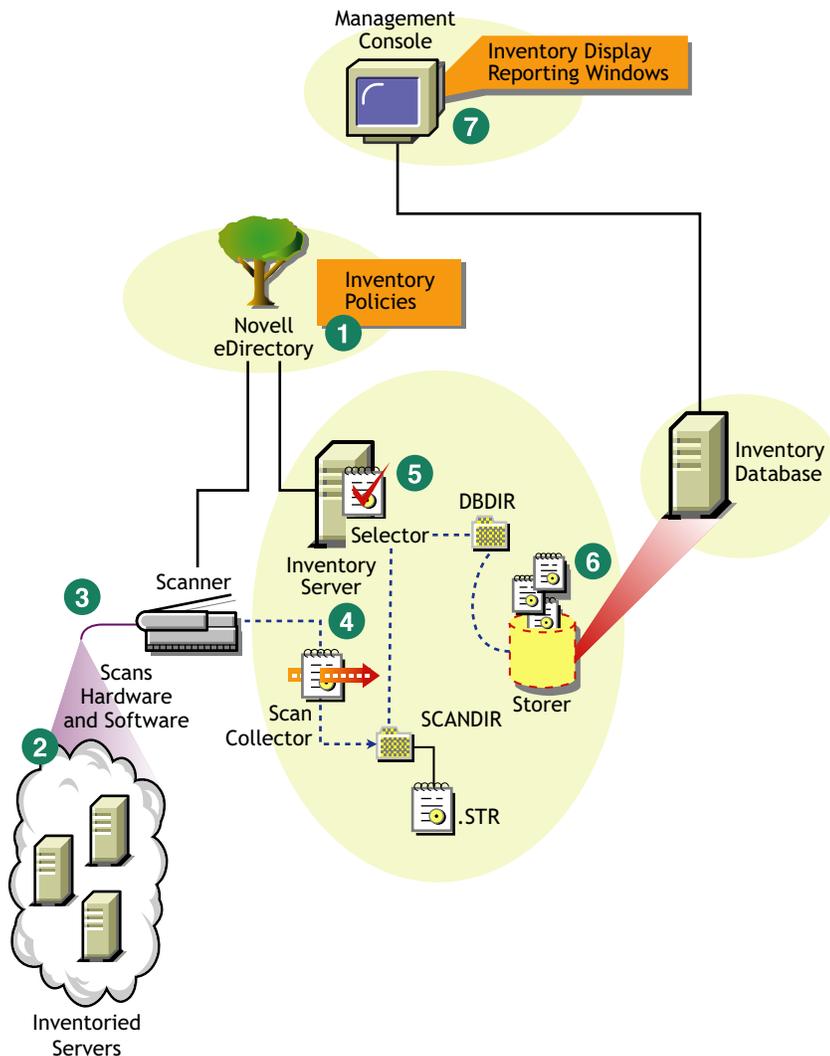
You must manually download the latest version of the dictionary from TID 10093255 in the [Novell Support Knowledgebase \(http://www.novell.com/support\)](#) and update the software dictionary.

NOTE: The dictionary is updated and published once every three months in this TID.

Understanding the Inventory Scanning Cycle in the Standalone Scenario

[Figure 13-2](#) depicts the scanning components and the inventory scanning cycle in the standalone scenario, which is explained below:

Figure 13-2 Inventory Scanning cycle in the Standalone Scenario



The inventory scanning cycle is as follows:

1. The inventory policies in the eDirectory define the inventory settings, such as the Inventory Service object name of the Inventory server to which the inventory information will be sent and the scanning time. These settings are customizable.
2. The scanner uses Policy and Distribution Services to read the inventory policies and collects the inventory information based on the policy settings. The Inventory scanner also checks whether an updated dictionary is available at its Inventory server and downloads the updated dictionary.
3. The scanner stores the inventory information locally on the inventoried server. This information is transferred to the Inventory server using the XML-RPC protocol.
4. The Scan Collector receives the inventory information using the XML-RPC protocol and stores the information in the scan directory at the Inventory server. The Scan Collector uses the ZENworks Web Server to process the XML-RPC requests.
5. The Selector validates the inventory information and places the information in the Inventory database.

6. The Storer updates the database with the inventory information.
7. The ZENworks administrator views the inventory information.

13.1.2 Advanced Deployment

- ♦ [“Deploying Inventory in a Single Site with More than 5,000 Inventoried Servers” on page 463](#)
- ♦ [“Deploying Inventory in Multiple or Enterprise Sites” on page 466](#)

Deploying Inventory in a Single Site with More than 5,000 Inventoried Servers

In this example scenario, the network consists of a single site with more than 5000 inventoried servers. The inventory configuration consists of two or more Standalone Servers; each server receiving scans from up to 5,000 inventoried servers. All the Standalone Servers store the inventory data to a single database.

The following sections contain detailed information to help you deploy Server Inventory on a single site:

1. [“Recommendations for Deployment” on page 463](#)
2. [“Installing Server Inventory” on page 464](#)
3. [“Understanding the Effects of Server Inventory Installation” on page 464](#)
4. [“Configuring the Required Policy” on page 464](#)
5. [“Starting the Inventory Service” on page 464](#)
6. [“Updating the Software Dictionary” on page 464](#)
7. [“Understanding the Inventory Scanning Cycle in the Standalone Scenario” on page 464](#)

Recommendations for Deployment

- ♦ The minimum base Inventory server configuration includes 512 MB RAM and a database cache of 128 MB.
- ♦ All inventoried servers should send the inventory information to the nearest Inventory server on the LAN; policies must be created based on this information.
- ♦ The transmission of inventory information to the Inventory server and storage of the inventory information into the Inventory database is an ongoing backend process that can take several hours or even more than a day.
- ♦ If many inventoried servers are attached to the same Inventory server, we recommend that you do not schedule the scan of all inventoried servers at the same time, because this stresses Novell eDirectory and the Inventory services.
- ♦ Ensure that the eDirectory time synchronization radius is set within 2 seconds.
- ♦ The optimal database cache size requirement for the server could vary because of the server environment. Determine the database cache size that needs to be set by trying a range of cache sizes in the runtime environment. The default Sybase database cache size is 128 MB. For more information about improving the database performance, see [Appendix J, “Performance Tips,” on page 717](#).

Installing Server Inventory

For detailed information on installing Server Inventory, see “[Policy-Enabled Server Management Installation](#)” in the *Novell ZENworks 7 Server Management Installation Guide*.

Understanding the Effects of Server Inventory Installation

For detailed information on the effects of Server Inventory installation, see [Section 13.1.3, “Understanding the Effects of Server Inventory Installation,”](#) on page 480.

Configuring the Required Policy

Configure the [Server Inventory Policy](#).

Starting the Inventory Service

After installing ZENworks 7 Server Management, the Inventory service is automatically started.

Updating the Software Dictionary

You can update the software dictionary in any one of the following ways:

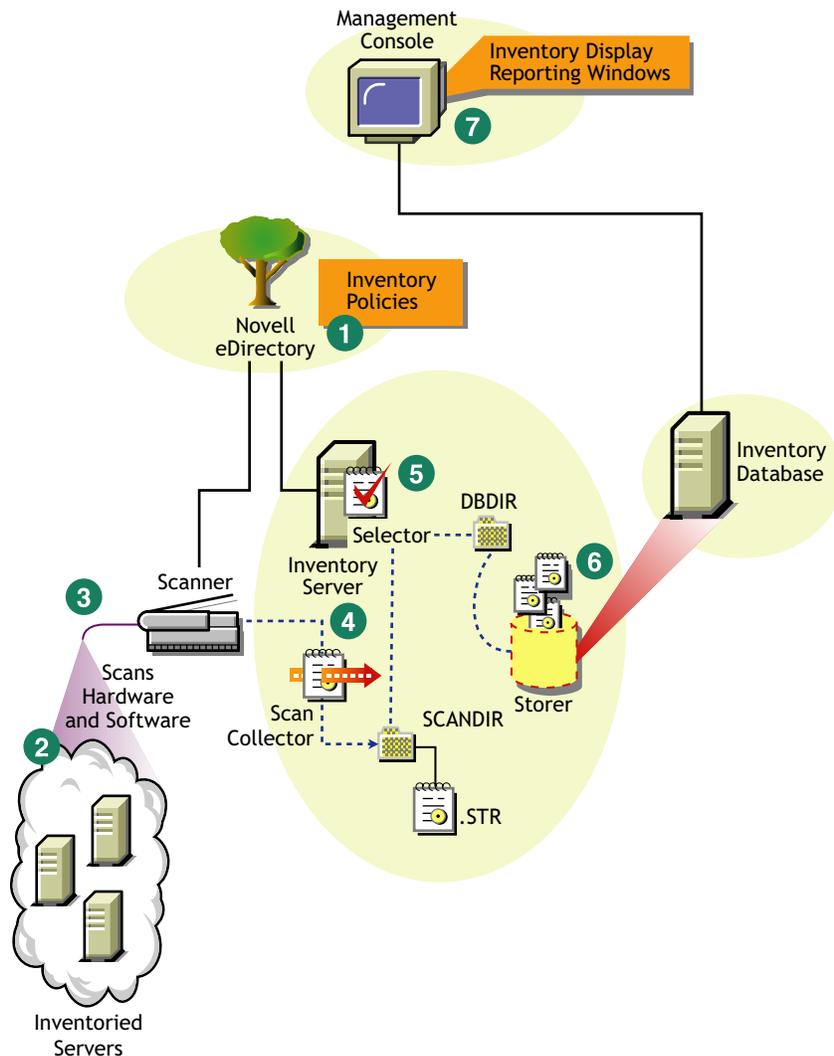
- ◆ On each Inventory server, manually download the latest version of the dictionary from TID 10093255 in the [Novell Support Knowledgebase \(http://www.novell.com/support\)](http://www.novell.com/support) and update the software dictionary.
- ◆ Manually download the latest version of the dictionary from TID 10093255 in the [Novell Support Knowledgebase \(http://www.novell.com/support\)](http://www.novell.com/support) on one of the Standalone Servers and automatically distribute the dictionary from this server to all other Standalone Servers in your setup by configuring the [Section 13.7, “Configuring the Dictionary Update Policy,”](#) on page 525. For more information, see [Section 13.8, “Setting Up Distribution of Dictionary,”](#) on page 526.

NOTE: The dictionary is updated and published once every three months in this TID.

Understanding the Inventory Scanning Cycle in the Standalone Scenario

[Figure 13-3](#) depicts the scanning components and the inventory scanning cycle in the standalone scenario, which is explained below:

Figure 13-3 Inventory Scanning Cycle in the Standalone scenario



The inventory scanning cycle is as follows:

1. The inventory policies in the eDirectory define the inventory settings, such as the Inventory Service object name of the Inventory server to which the inventory information will be sent and the scanning time. These settings are customizable.
2. The scanner uses Policy and Distribution Services to read the inventory policies and collects the inventory information based on the policy settings. The Inventory scanner also checks whether an updated dictionary is available at its Inventory server and downloads the updated dictionary.
3. The scanner stores the inventory information locally on the inventoried server. This information is transferred to the Inventory server using the XML-RPC protocol.
4. The Scan Collector receives the inventory information using the XML-RPC protocol and stores the information in the scan directory at the Inventory server. The Scan Collector uses the ZENworks Web Server to process the XML-RPC requests.
5. The Selector validates the inventory information and places the information in the Inventory database.

6. The Storer updates the database with the inventory information.
7. The ZENworks administrator views the inventory information.

Deploying Inventory in Multiple or Enterprise Sites

The following sections contain detailed information to help you deploy Server Inventory in multiple or enterprise sites:

1. [“Designing the Inventory Tree” on page 466](#)
2. [“Deployment Options for Inventory Server and Inventory Database” on page 471](#)
3. [“Recommendations for Deployment” on page 476](#)
4. [“Installing Server Inventory” on page 476](#)
5. [“Understanding the Effects of Server Inventory Installation” on page 477](#)
6. [“Configuring the Required Policies” on page 477](#)
7. [“Starting the Inventory Service” on page 478](#)
8. [“Updating the Software Dictionary” on page 478](#)
9. [“Understanding Rolling Up Inventory Information Across Servers” on page 478](#)

Designing the Inventory Tree

In an enterprise or multiple site, complete the following tasks to design the inventory tree:

- ◆ [“1. List the sites in the enterprise” on page 466](#)
- ◆ [“2. What is the ideal place for the Root Server?” on page 467](#)
- ◆ [“3. Is any other database needed?” on page 467](#)
[“Optional step: If another database is needed” on page 468](#)
- ◆ [“4. Identify the route for Inventory information” on page 468](#)
- ◆ [“5. Identify servers on each site to act as Inventory and Database Servers” on page 468](#)
- ◆ [“6. Create the tree of servers for enterprise Inventory collection” on page 469](#)
- ◆ [“7. Create an implementation plan” on page 470](#)
- ◆ [“8. Start the actual deployment” on page 470](#)

1. List the sites in the enterprise

Describe the entire network of your enterprise.

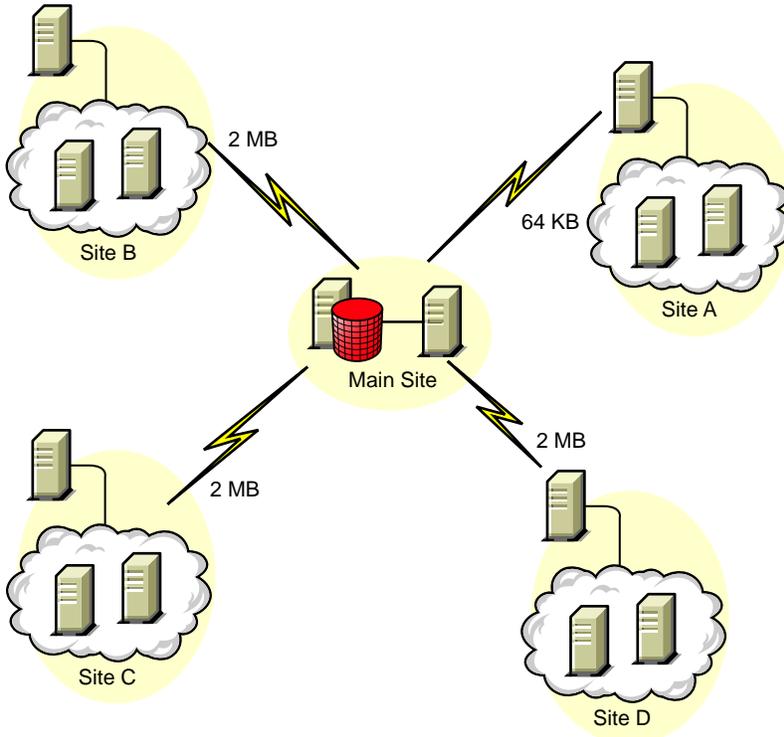
- ◆ List the various sites in your enterprise.
- ◆ List the physical links between the various sites.
- ◆ Identify the type of links in terms of bandwidth and reliability.

Figure 13-4 illustrates the network organization of an enterprise with servers in different locations:

Figure 13-4 Network organization of an enterprise

Network Configuration of My Company

No. of NetWare Servers = 2
No. of Windows NT Servers = 5



This illustration depicts four sites (Site A, Site B, Site C, and Site D) connected to a central site. It depicts the physical links between the sites and the type of links in terms of bandwidth.

2. What is the ideal place for the Root Server?

The Root Server in the inventory tree is the highest-level server. Necessarily, an Inventory database is attached to the Root Server.

The inventory information available from the Inventory database of the Root Server consists of all information from lower-level sites on the network and from the Root Server site.

Factors that you must consider include:

- ◆ There must be high-speed links between the Root Server and the management console.
- ◆ We recommend that there should be high-speed links between the site having the Root Server and the sites having the lower-level Inventory servers.
- ◆ Using the management console, the administrator can collect the inventory information from any of the sites connected on high-speed links from the Root Server, or from the Root Server level site.
- ◆ A database server of suitable configuration should be provided for the Inventory server.

3. Is any other database needed?

In addition to the database at the Root Server, you can maintain database servers at different sites.

You might want to maintain additional databases if there are sites or subtrees that are managed for inventory at different locations, and these sites are connected to the network over a slow link.

You should also determine if there are specific reasons to have a separate database for a single site or a set of sites. There might be some organizational needs for your enterprise to have the database server on different sites, even if there is no product deployment need to have any other database.

NOTE: For a majority of enterprises, there could be no need to have any other database besides the enterprise-wide single database.

Optional step: If another database is needed

- ◆ If you decide to have additional database servers, identify the sites that need a database. Additionally, you need to examine whether the database caters to the local site or a site with many subsites. Also, identify the sites that require information in each Inventory database.
- ◆ All the sites served by a single database should typically access this database instead of the database at the Root Server for inventory management. This reduces the load on the database at Root Server.
- ◆ Database administrators should be available for these sites.

4. Identify the route for Inventory information

Identify the routes for inventory information for all Inventory servers to the nearest database.

To devise a route plan:

- ◆ Each route can have an Intermediate Server at a staging site. The Intermediate Server receives and transmits the information to the next destination. These are application-layer-level routes for inventory information. There can be various network-layer-level routes between two adjacent servers, which is determined and managed by the routers in the network.
- ◆ The route provides information indicating how inventory information travels from a particular site to its final destination, which is the database at the Root Server.
- ◆ There can be multiple routes. Choose the fastest and most reliable route. To determine the route, consider the physical network links.
- ◆ Routes identified and made operational can be changed later, although there might be some cost in terms of management and traffic generation. If there is no intermediate database involved, you can change the route by changing the eDirectory-based policy.
- ◆ Put Intermediate Servers on sites where the link parameters change substantially. Criteria to consider are difference in bandwidth, difference in reliability of the links, and the need for roll up of inventory information.
- ◆ Availability of Inventory servers on the intermediate site for staging the inventory information should be considered while deciding the sites for Intermediate Servers. Ensure that there is enough disk space on these servers to store all the inventory information on the disk until the Sender sends it to the next destination.

5. Identify servers on each site to act as Inventory and Database Servers

In ZENworks 7 Server Management, you choose the role for each Inventory server. For more information, see [Section 12.4, “Understanding the Inventory Server Roles,” on page 448](#).

The number of inventoried servers attached to an Inventory server also determines the load. [Table 13-1](#) lists the disk space requirements for the server:

Table 13-1 *Disk Requirements for a ZENworks Inventory Server*

| Server Type | Disk Space Requirements |
|---|--|
| Leaf Server | $(n1 \times s) + (n1 \times z)$ |
| Leaf Server with Database | $(n1 \times s \times 2) + \{(n1 \times dbg)\}$ |
| Intermediate Server | $n2 \times z$ |
| Intermediate Server with Database | $(n2 \times z) + (n2 \times s) + \{(n2 \times dbg)\}$ |
| Intermediate Server with Inventoried Servers | $(n1 \times s \times 2) + (n2 \times z)$ |
| Intermediate Server with Database and Inventoried Servers | $(n1 \times s \times 2) + (n2 \times z) + (n2 \times s) + \{(n1 \times dbg) + (n2 \times dbg)\}$ |
| Root Server | $(n2 \times z) + (n2 \times s) + \{(n2 \times dbg)\}$ |
| Root Server with Inventoried Servers | $(n1 \times s \times 2) + (n2 \times z) + (n2 \times s) + \{(n1 \times dbg) + (n2 \times dbg)\}$ |
| Standalone Server | $(n1 \times s \times 1) + \{(n1 \times dbg)\}$ |

In the table, $n1$ is the number of inventoried servers attached to the server.

s is the size of the scan data files. This file size varies depending on the information collected. Calculate 400 KB scan information from each inventoried server to calculate the load.

dbg is the storage space of the inventory information in the database.

$n2$ is the number of inventoried servers rolled up to the Inventory server.

z is the size of the compressed scan data file per inventoried server. The average compression ratio is 80-90% of the STR file size.

{ } denotes the disk space of the database server, depending on whether the database is on the same Inventory server or if it is connected to the Inventory server. If the database is on the same Inventory server, calculate the total disk space including the database space for the Inventory server. For example, if the Leaf Server with Database has the Inventory database on the same server, calculate the requirements for storage of inventory information, including the database disk space.

6. Create the tree of servers for enterprise Inventory collection

Ensure that the inventory tree you design follows these guidelines:

- ◆ The root of the tree is the Root Server.
- ◆ At least one Inventory server per site is recommended.
- ◆ Each site has inventoried servers to be scanned.
- ◆ Optionally, there are databases and Intermediate Servers on different sites.

7. Create an implementation plan

After you design the inventory tree, you should develop an implementation plan to cover the phased deployment plan for the network. Use the top-down deployment of the Server Inventory installation. Always begin the installation at the topmost level server (Root Server) and proceed with the next lower-level servers.

8. Start the actual deployment

After your implementation plan is finalized, start the actual deployment according to the plan.

Follow these steps:

1. Install the Inventory servers on the sites. For more information, see [“Installing Server Inventory” on page 476](#).
2. Create and configure the policies applicable to Inventory server and inventoried servers. For more information, see [“Configuring the Required Policies” on page 477](#).

Adding a Database Server to an Existing Inventory Setup

If you have already configured the servers for inventory setup, and you need to add another database server, follow these instructions:

- 1** Run the installation program to install the Inventory database on the server.

The installation program installs the Sybase database. If you are maintaining the database in Oracle, make sure that the Oracle database exists. For more information, see [“Setting Up the Oracle Inventory Database” on page 500](#). If you are maintaining the database in MS SQL, make sure that the MS SQL database exists. For more information, see [“Setting Up the MS SQL Server 2000 or MS SQL Server 2005 Inventory Database” on page 510](#).

- 2** Shut down the Inventory services. For more information, see [“Stopping the Inventory Service on a NetWare Inventory Server” on page 482](#).
- 3** Based on the database you select, make sure that you configure the database. For more information, see [Section 13.4, “Configuring the Database Location Policy,” on page 521](#).
- 4** If you want to attach a new database to an existing Inventory server that does not have a database attached, you must change the role of the Inventory server in its Inventory Service object (Inventory Service_ *server_name*). For example, if you attach an Inventory database to an existing Leaf Server, you must change the server role from Leaf Server to Leaf Server with Database in the Inventory Service object of the Leaf Server.

If you want to attach an Inventory database to a freshly installed Inventory server, you must choose an appropriate server role for the Inventory server.

To change the role of an Inventory server:

- 4a** In ConsoleOne, right-click the Inventory Service object, click *Properties*, then click the *Inventory Service Object Properties* tab.
 - 4a1** Choose the new role of the Inventory Service object, then click *Apply*.

Follow the actions that you need to change the role. For more information, see [Section 13.1.5, “Changing the Role of the Inventory Server,” on page 483](#).
- 5** Make sure that you enforce Full Scan for the Inventory Service object:
 - 5a** In ConsoleOne, right-click the Inventory Service object, click *Properties*, then click the *Inventory Service Object Properties* tab.

- 5b** Select the *Enable Scan* option, then click *OK*.
- 6** Bring up the Inventory service. For more information, see [“Starting the Inventory Service on a NetWare Inventory Server” on page 482](#).

Deployment Options for Inventory Server and Inventory Database

The following sections cover these scenarios:

- ◆ [“Scenario 1: Inventory Deployment without Intermediate Servers in a WAN” on page 471](#)
- ◆ [“Scenario 2: Inventory Deployment with Intermediate Servers in a WAN” on page 472](#)
- ◆ [“Scenario 3: Roll Up of the Inventory Information Across eDirectory Trees” on page 474](#)
- ◆ [“Scenario 4: Merging eDirectory Trees” on page 475](#)
- ◆ [“Scenario 5: Deploying Inventory Server Across a Firewall” on page 475](#)

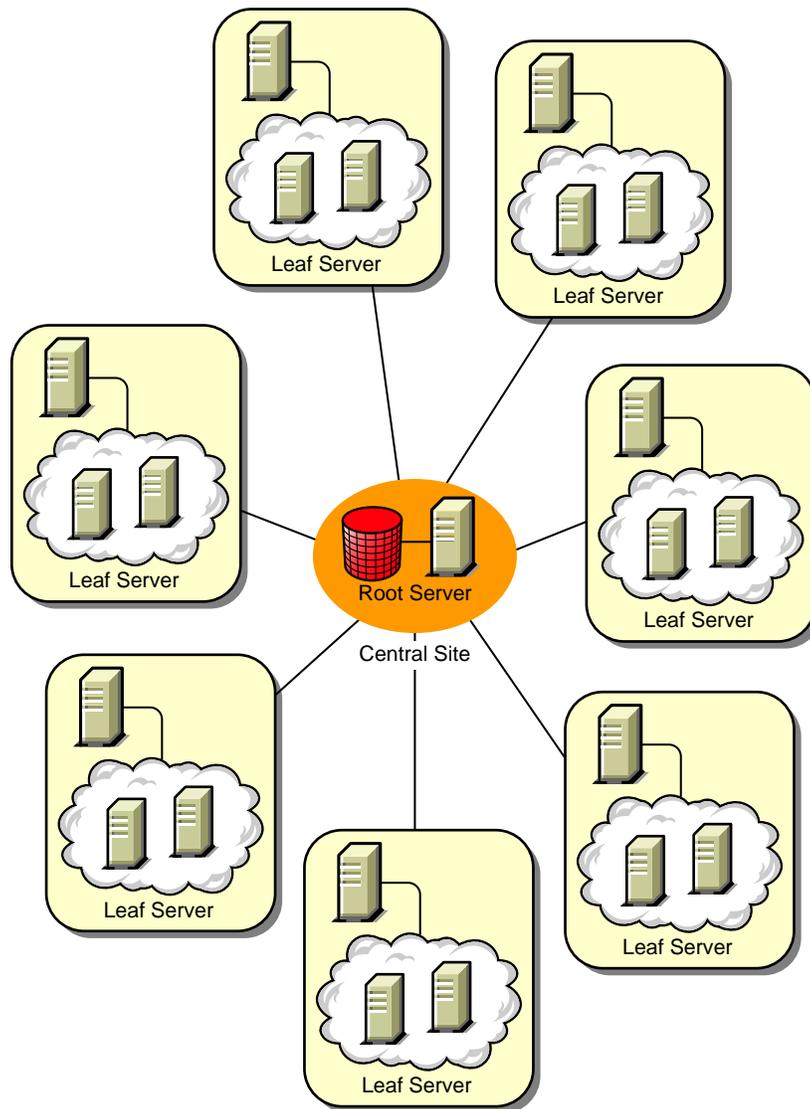
Scenario 1: Inventory Deployment without Intermediate Servers in a WAN

In this scenario, the network consists of many remote sites connected to a Central Site over a WAN. Each remote site has a Leaf Server that collects inventory information from inventoried servers located in the same site, and rolls up the inventory information to the Root Server located at the central site. The remote sites are administered from the Central Site because the Leaf Servers do have Inventory database attached to it.

TIP: To locally administer the remote sites, you must have the Inventory database attached to Leaf Servers and change the role of the Inventory server to Leaf Server with Database. For more information on how to change the role of an Inventory server, see [Section 13.1.5, “Changing the Role of the Inventory Server,” on page 483](#).

This scenario is illustrated in [Figure 13-5](#):

Figure 13-5 Inventory Deployment without Intermediate Servers in a WAN

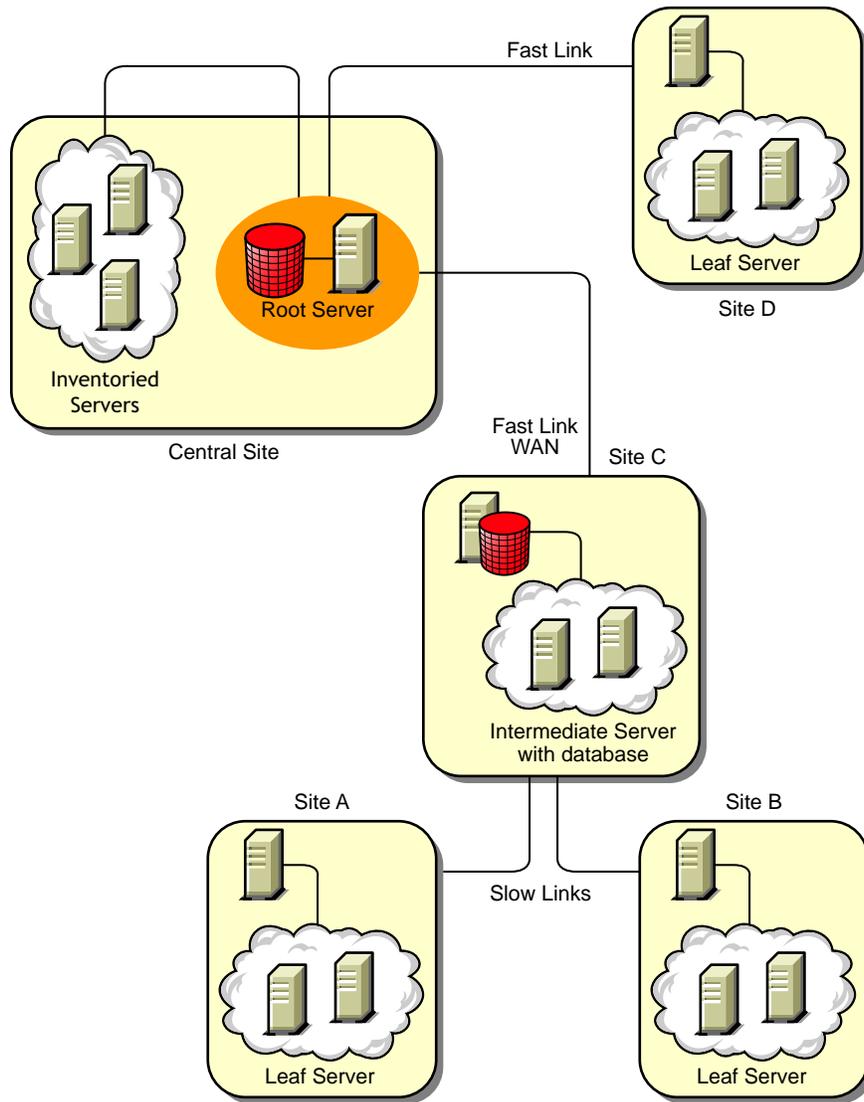


Scenario 2: Inventory Deployment with Intermediate Servers in a WAN

In this scenario, the network consists of four remote sites (A, B, C, and D), and a Central Site. Sites A and B are connected to Site C over slow links and are not directly connected to the Central Site. Site C is connected to the Central Site over a fast WAN link. Site D is directly connected to the Central Site over a fast link. Sites A, B and C are administered at Site C.

This scenario is illustrated in [Figure 13-6](#):

Figure 13-6 Inventory Deployment with Intermediate Servers in a WAN



To administer the enterprise from the Central Site, and also administer Sites A and B from Site C, do the following:

1. Install Leaf Servers at Sites A, B, and D.
2. Install Intermediate Servers with Database at Site C.
3. Configure Leaf Servers at Sites A and B to roll up the inventory information to the Intermediate Server with Database at Site C.
4. Configure the Intermediate Server with Database at Site C to roll up the inventory information to a Root Server at the Central Site.
5. Configure the Leaf Server at Site D to roll up the inventory information to a Root Server at the Central Site.

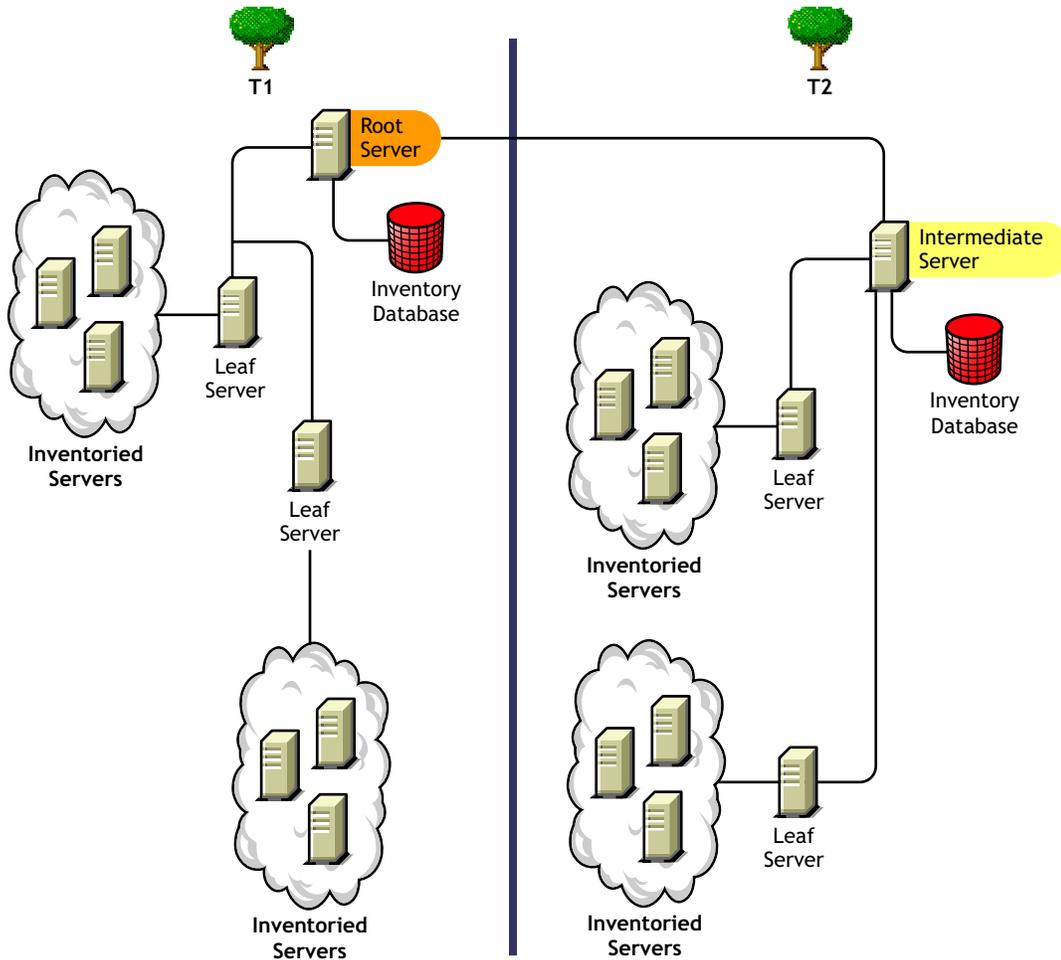
Scenario 3: Roll Up of the Inventory Information Across eDirectory Trees

In this configuration, you can deploy any of the previous scenarios. The highest-level Inventory server of one eDirectory tree rolls up the inventory information to an Inventory server located on the other eDirectory tree.

In this configuration, you must install the Distributor on each eDirectory tree for the policies to be distributed.

Figure 13-7 depicts a sample scenario where you can deploy this inventory configuration:

Figure 13-7 Roll Up of the Inventory Information across eDirectory Trees



There are two organizations: A and B. Each organization has its own eDirectory tree and inventory tree. Organization A has two Leaf Servers and a Root Server in its inventory tree. Organization B also has two Leaf Servers and a Root Server in its inventory tree. A decision is taken to merge both the organizations and both the inventory trees but to retain the eDirectory trees. After the merger, the role of the Root Server on the eDirectory tree T2 is changed to Intermediate Server with Database and the inventory information is rolled up from the Intermediate Server to the Root Server residing on the eDirectory tree T1.

Scenario 4: Merging eDirectory Trees

In this configuration, you can merge the inventory trees and the eDirectory trees. After you merge the eDirectory trees, you must manually change the eDirectory tree name and (optionally) the Inventory Service DN in the `inventory_server_installation_drive_or_volume\zenworks\inv\server\wminv\properties\config.properties` file before starting the Inventory service. For more information on merging the eDirectory trees, see the [Novell eDirectory documentation Web site \(http://www.novell.com/documentation\)](http://www.novell.com/documentation).

To merge the inventory trees, you must change the role of the Root Server of one inventory tree to roll up to an Inventory server in the other inventory tree.

To change the eDirectory tree name and the DN of an Inventory server, edit the following entries of the `config.properties` file:

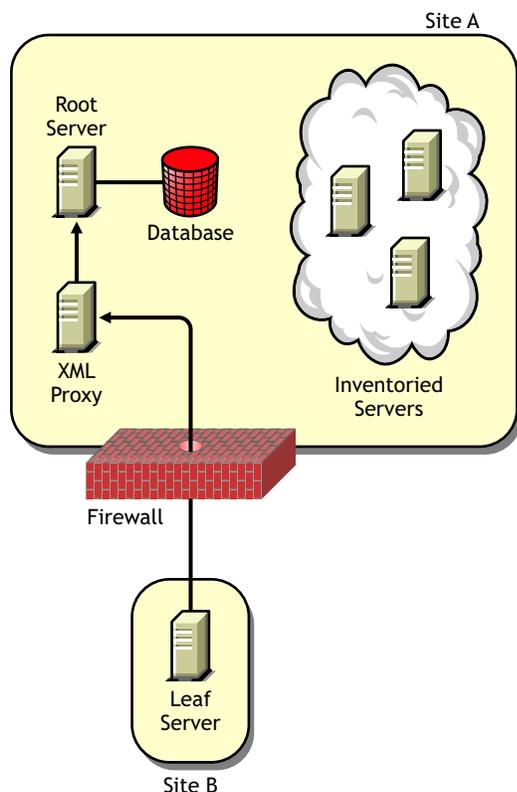
```
NDSTree=Target_eDirectory_tree_name
InventoryServiceDN=New_DN_of_the_Inventory_server
```

Scenario 5: Deploying Inventory Server Across a Firewall

There are two sites; Site A and Site B connected through a WAN link. The Inventory server of Site A rolls up to an Inventory server in Site B. All communication from Site A to Site B flows through the firewall at Site B.

Figure 13-8 depicts a sample scenario where you can deploy this inventory configuration:

Figure 13-8 Deploying Inventory Server across a Firewall



To enable the roll-up:

- ◆ Install an XML proxy at Site A. For more information about installing the proxy, see “[Policy-Enabled Server Management Installation](#)” in the *Novell ZENworks 7 Server Management Installation Guide*.
- ◆ You must have at least one XML proxy/site installed. One proxy server can handle requests for multiple Inventory servers.
- ◆ You can configure the port that the proxy listens to during the ZENworks 7 Server Management installation. For more information, see “[Policy-Enabled Server Management Installation](#)” in the *Novell ZENworks 7 Server Management Installation Guide*.

You must allow requests to the proxy server on this port at the firewall. You can configure the XML proxy to listen to standard ports allowed by your firewall.

The XML proxy does not support any commercial Web server. You must ensure that the port number assigned to the XML proxy is not used by any other service on the same server.

You must also configure the Roll-Up policy with the XML proxy server's address and port number.

Recommendations for Deployment

- ◆ When you configure the inventory scanning of inventoried servers, we recommend staggering the inventory scanning to scan at different times or to scan a group of inventoried servers at a time.
- ◆ If many inventoried servers are attached to the same inventory server, we recommend that you do not schedule the scan of all inventoried servers at the same time, because this stresses Novell eDirectory and the Inventory service.
- ◆ You can attach inventoried servers to the server as determined by the number of connections supported by NetWare, Linux, or Windows servers up to a maximum of 5,000 inventoried servers.
- ◆ When you schedule the roll-up of information in the Inventory policies, we recommend the roll-up frequency should be at least one day. If the roll-up of inventory information is scheduled too frequently, for example less than one hour, there may be some performance degradation of the inventory server.
- ◆ Use top-down deployment for Inventory installation. Always begin the installation at the topmost level server and proceed with the next lower-level servers. For example, in an inventory setup with a Root Server and a Leaf Server, complete the inventory installation at the Root Server, and then run the installation for the Leaf Server.
- ◆ If an Inventory server must receive server Inventory scans either directly from the Inventory servers or through roll-up, you must install ZENworks 7 Server Management on this server.
- ◆ We recommend that you configure DNS for your Inventory and database servers. If you have not configured DNS, choose the IP address in the Roll-Up and Database Location policies. Scheduling the frequency of information gathering and roll-up must be fine-tuned based on the Root Server. Make sure that the Root Server is able to handle the load of the .STR files.

Installing Server Inventory

For detailed information on installing Server Inventory, see “[Policy-Enabled Server Management Installation](#)” in the *Novell ZENworks 7 Server Management Installation Guide*.

Understanding the Effects of Server Inventory Installation

For detailed information on the effects of Server Inventory installation, see [Section 13.1.3, “Understanding the Effects of Server Inventory Installation,”](#) on page 480.

Configuring the Required Policies

[Table 13-2](#) lists the actions that you should follow to set up the server for Server Inventory.

Table 13-2 Policies Required to set up an Inventory Server

| To set up this type of server: | Do this: |
|--|--|
| Standalone Server | <ol style="list-style-type: none">1. Follow the steps in Section 13.4, “Configuring the Database Location Policy,” on page 5212. Follow the steps in Section 13.5, “Configuring the Server Inventory Policy,” on page 522 |
| Root Server | <ol style="list-style-type: none">1. Follow the steps in Section 13.3, “Configuring the Inventory Service Object,” on page 5202. Follow the steps in Section 13.4, “Configuring the Database Location Policy,” on page 521 |
| Root Server with Inventoried Servers | <ol style="list-style-type: none">1. Follow the steps in Section 13.3, “Configuring the Inventory Service Object,” on page 520.2. Follow the steps in Section 13.5, “Configuring the Server Inventory Policy,” on page 522.3. Follow the steps in Section 13.4, “Configuring the Database Location Policy,” on page 521. |
| Intermediate Server | <ol style="list-style-type: none">1. Follow the steps in Section 13.3, “Configuring the Inventory Service Object,” on page 520.2. Follow the steps in Section 13.6, “Configuring the Roll-Up Policy,” on page 524. |
| Intermediate Server with Database | <ol style="list-style-type: none">1. Follow the steps in Section 13.3, “Configuring the Inventory Service Object,” on page 520.2. Follow the steps in Section 13.6, “Configuring the Roll-Up Policy,” on page 524.3. Follow the steps in Section 13.4, “Configuring the Database Location Policy,” on page 521. |
| Intermediate Server with Inventoried Servers | <ol style="list-style-type: none">1. Follow the steps in Section 13.3, “Configuring the Inventory Service Object,” on page 520.2. Follow the steps in Section 13.5, “Configuring the Server Inventory Policy,” on page 522.3. Follow the steps in Section 13.6, “Configuring the Roll-Up Policy,” on page 524. |

| To set up this type of server: | Do this: |
|---|---|
| Intermediate Server with Database and Inventoried Servers | <ol style="list-style-type: none"> 1. Follow the steps in Section 13.3, "Configuring the Inventory Service Object," on page 520. 2. Follow the steps in Section 13.5, "Configuring the Server Inventory Policy," on page 522. 3. Follow the steps in Section 13.6, "Configuring the Roll-Up Policy," on page 524. 4. Follow the steps in Section 13.4, "Configuring the Database Location Policy," on page 521. |
| Leaf Server | <ol style="list-style-type: none"> 1. Follow the steps in Section 13.3, "Configuring the Inventory Service Object," on page 520. 2. Follow the steps in Section 13.5, "Configuring the Server Inventory Policy," on page 522. 3. Follow the steps in Section 13.6, "Configuring the Roll-Up Policy," on page 524. |
| Leaf Server with Database | <ol style="list-style-type: none"> 1. Follow the steps in Section 13.3, "Configuring the Inventory Service Object," on page 520. 2. Follow the steps in Section 13.5, "Configuring the Server Inventory Policy," on page 522. 3. Follow the steps in Section 13.6, "Configuring the Roll-Up Policy," on page 524. 4. Follow the steps in Section 13.4, "Configuring the Database Location Policy," on page 521. |

Starting the Inventory Service

After installing ZENworks 7 Server Management, the Inventory service is automatically started only if you have configured the Inventory Standalone Configuration settings during the installation.

To manually start the Inventory service, see ["Starting the Inventory Service on a NetWare Inventory Server"](#) on page 482.

Updating the Software Dictionary

You can update the software dictionary in any one of the following ways:

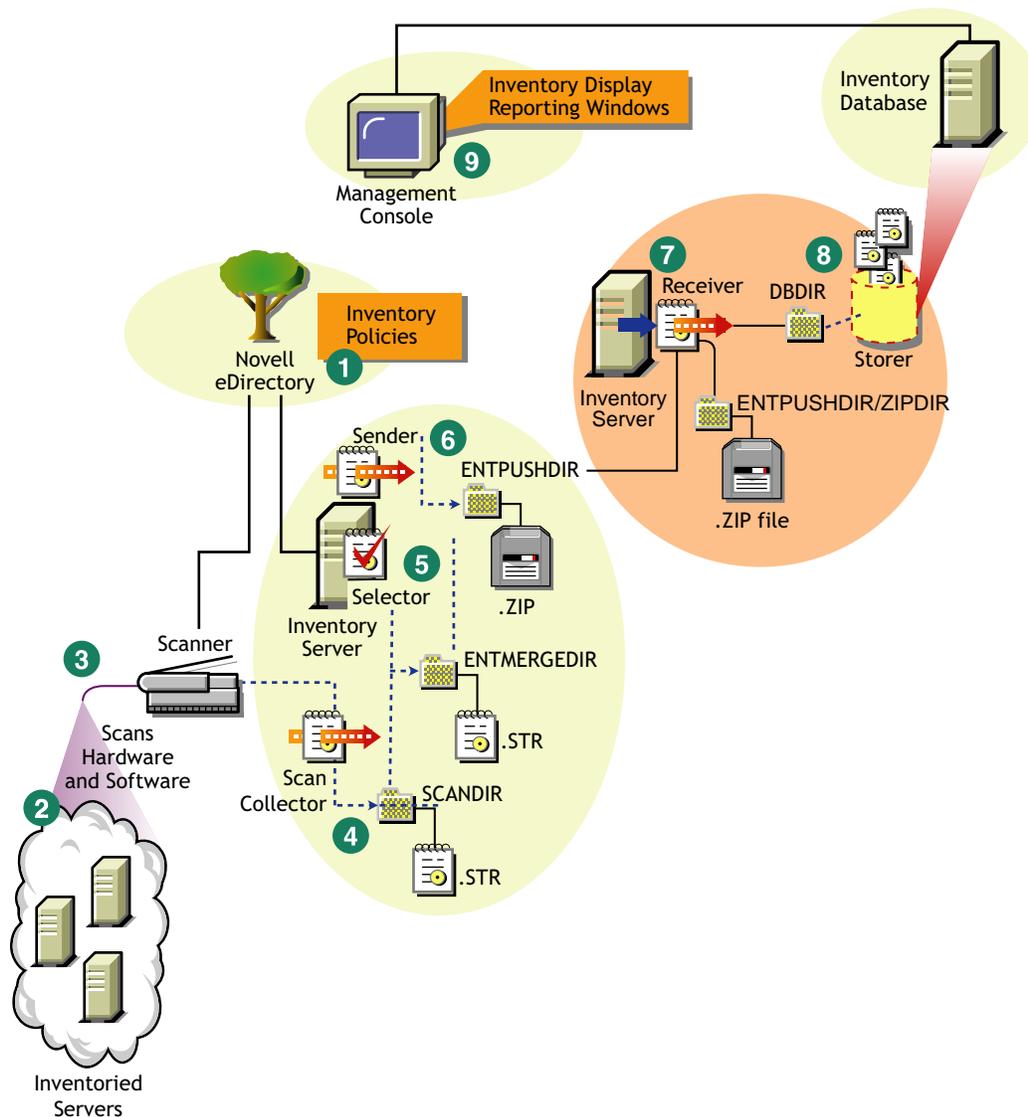
- ◆ On each Inventory server, manually download the latest version of the dictionary from TID 10093255 in the [Novell Support Knowledgebase \(http://www.novell.com/support\)](http://www.novell.com/support) and update the software dictionary.
- ◆ Manually download the latest version of the dictionary from TID 10093255 in the [Novell Support Knowledgebase \(http://www.novell.com/support\)](http://www.novell.com/support) on an Inventory server (preferably, the Root Server) and automatically distribute the dictionary to all servers in your setup by configuring the [Section 13.7, "Configuring the Dictionary Update Policy,"](#) on page 525. For more information, see [Section 13.8, "Setting Up Distribution of Dictionary,"](#) on page 526.

NOTE: The dictionary is updated and published once every three months in this TID.

Understanding Rolling Up Inventory Information Across Servers

[Figure 13-9](#) depicts rolling up the inventory information across servers, which is explained below:

Figure 13-9 Inventory scanning cycle in the Roll-Up scenario



If the inventory deployment rolls up inventory information across servers, the process of scanning is as follows:

1. The inventory policies in eDirectory define the inventory settings, such as the Inventory Service object name of the Inventory server to which the inventory information will be sent, scanning time, and the software rules for software scan. These settings are customizable.
2. The Scanner uses Policy and Distribution Services to read the inventory policies and collects the inventory information based on the policy settings. The Inventory scanner also checks whether an updated dictionary is available at its Inventory server and downloads the updated dictionary.
3. The Scanner stores the inventory information locally on the inventoried server. This information is transferred to the Inventory server using the XML-RPC protocol.

4. The Scan Collector receives the inventory information using the XML-RPC protocol and stores the `.str` file in the scan directory at the Inventory server. The Scan Collector uses the ZENworks Web Server to process the XML-RPC requests.
5. The Selector validates the inventory information and places it in the `enterprise merge` directory for roll-up of inventory information. If there is a database attached, the Selector also places the files in the database directory.
6. The Sender on the Inventory server has a Roll-Up policy to identify the Inventory server to which it will transmit the inventory information and the Roll-Up schedule specifies the time for roll-up of information. The Sender compresses the `.str` files as a `.zip` file and places the `.zip` file in the `enterprise push` directory. The Sender then sends the `.zip` file to the Receiver on the next-level Inventory server.
7. The Receiver on the next-level Inventory server receives the `.zip` file.

NOTE: The next-level Inventory server can be located on the same eDirectory tree or on a different eDirectory tree.

On the Intermediate Server, the Receiver copies the file in the `enterprise push` directory (`entpushdir`). On the Intermediate Server with Database, or the Intermediate Server with Database and Inventoried Servers, the Receiver places the file in `\entpushdir` and places the file to the database directory (`dbdir`).

On the Root Server, or the Root Server with Inventoried Servers, the Receiver copies the file to the `\dbdir` only.

8. The Storer extracts the `.zip` file containing the `.str` files to a temp directory (`\dbdir\temp`) and updates the database with the inventory information of the inventoried server `.str` file.
9. The ZENworks administrator views the inventory information, and queries the database in ConsoleOne.

13.1.3 Understanding the Effects of Server Inventory Installation

On the Inventory server, the ZENworks 7 Server Inventory installation program does the following:

- ◆ “On NetWare Inventory Servers” on page 480
- ◆ “On Windows Inventory Servers” on page 481
- ◆ “On Database Servers” on page 481

On NetWare Inventory Servers

- ◆ Copies the inventory related files to the `installation_volume`.
- ◆ Copies the Server Inventory snap-ins to the ConsoleOne directory.
- ◆ Creates an Inventory Service object in eDirectory for each server on which the Inventory server is installed. This object is populated with the following attributes: `zeninvRole` (role of the server), `zeninvScanFilePath` (path to `\scandir` directory) `zeninvHostServer` (DN of the server on which Inventory server is installed), and `zeninvDictionarypath` (Path to the dictionary directory).
- ◆ If the Inventory Service object already exists, the object is validated and re-created if it is invalid.

- ◆ During installation, the Inventory Service object is made a trustee of the NCP™ server with Compare and Read rights.
- ◆ Assigns the Inventory Service object as trustee to itself.
- ◆ Creates the scan directory with the subdirectories in the specified volume on the Inventory server. [Root] is granted the Create rights to this directory.
- ◆ Creates a dictionary directory (`dictdir`), and copies the files of general dictionary and private dictionary. [Root] is granted the Read and Write rights to this directory.
- ◆ Creates the `zenworks.properties` file in `sys:\system`. This file contains the installation path of the Inventory server and the ZENworks Web Server.
- ◆ During the Server Inventory installation, if you have configured Inventory Standalone Configuration settings, then the Inventory Service Manager is automatically started.
- ◆ Installs the ZENworks Web Server on the Inventory server, if not installed previously.
- ◆ If Server Inventory is reinstalled in the same directory as the previous installation, the `config.properties` and `directory.properties` files are backed up and re-created.

On Windows Inventory Servers

- ◆ Copies the inventory related files to the *installation_directory*.
- ◆ Copies the Server Inventory snap-in component to the ConsoleOne directory.
- ◆ Creates the scan directory with the subdirectories in the specified volume on the Inventory server, and creates a share with Create rights to this directory for all users.
- ◆ Creates a dictionary directory (`dictdir`), copies the files of general dictionary and private dictionary, and grants Read and Write rights to this directory for all users.
- ◆ Creates an Inventory Service object in eDirectory for each server on which the Inventory server is installed. The following attributes are populated: `zeninvRole` (role of the server), `zeninvScanFilePath` (path to `\scandir`), `zeninvHostServer` (DN of the server on which Inventory is installed), and `zeninvDictionarypath` (path to the dictionary directory).
- ◆ If the Inventory Service object already exists, the object is validated and re-created if it is invalid.
- ◆ During installation, the Inventory Service object is made a trustee of the NCP server with Compare and Read rights.
- ◆ The installation program assigns the Inventory Service object as trustee to itself.
- ◆ The Inventory Service Manager is created as a service.
- ◆ Edits the Registry settings to add the installation path of the Inventory server and the ZENworks Web Server.
- ◆ On the Inventory server, the ZENworks Service Management is created as a service.
- ◆ If Server Inventory is reinstalled in the same directory as the previous installation directory, the `config.properties` and `directory.properties` files are backed up and re-created.

On Database Servers

- ◆ Installs the Sybase database on the server you specify.
- ◆ If the database server is installed in the previous installation directory, the database files are re-created if they were found invalid or non-existing.
- ◆ If Sybase is already installed, only the database files are copied.

- ◆ On NetWare, the `mgmtdb.db` entries are added to the `sys:\system\mgmt dbs.ncf` file. On Windows, the `mgmtdb.db` entries are added to the registry.
- ◆ Creates a database object (Inventory database `_server_name`) for Sybase and configures the properties of the object.
- ◆ At server startup time, the database is loaded.

13.1.4 Starting and Stopping the Inventory Service

The section provides information on:

- ◆ [“Starting the Inventory Service on a NetWare Inventory Server” on page 482](#)
- ◆ [“Stopping the Inventory Service on a NetWare Inventory Server” on page 482](#)
- ◆ [“Starting the Inventory Service on a Windows Inventory Server” on page 482](#)
- ◆ [“Stopping the Inventory Service on a Windows Inventory Server” on page 483](#)

For more information about the various Inventory services, see [Section 14.1, “Understanding the Inventory Service Manager,” on page 529](#).

Starting the Inventory Service on a NetWare Inventory Server

Before you start the Inventory service, make sure that the Inventory database is up and running. The Inventory database is automatically started after the installation.

To start the Inventory service on the NetWare Inventory server, enter `startinv` at the server console prompt.

To start an Inventory service, enter `startser inventory_service_name` at the server console prompt.

After starting the Inventory service, make sure that the Inventory services are up and running. To list all services, enter `listser *` at the server console prompt. To list an Inventory service, enter `listser inventory_service_name` at the server console prompt.

If the services are not up and running, check the Server Status log. For more information on the Server Status log, see [Section 18.3, “Viewing the Status of Inventory Components on an Inventory Server,” on page 712](#).

Stopping the Inventory Service on a NetWare Inventory Server

To stop an Inventory service, enter `stopser inventory_service_name` at the server console prompt.

To stop all the Inventory services, enter `stopser *` at the server console prompt.

Starting the Inventory Service on a Windows Inventory Server

Before you start the Inventory service, make sure that the ZENworks Server Management components and the Inventory database are up and running. The Inventory database is automatically started after the installation.

To start the Inventory services on the Windows 2000/2003 Inventory server:

- 1 In the Control Panel, double-click *Administrative Tools*.

- 2 Double-click *Services*.
- 3 Select *Novell Inventory Service*, then click *Start*.

To start an Inventory service from the console prompt:

- 1 Go to the `installation_directory\inv\server\wminv\bin` directory.
- 2 At the prompt, enter `startser inventory_service_name`.

After starting the Inventory service, make sure that the Inventory services are up and running. To list all services, enter `listser "*"` at the server console prompt. To list an Inventory service from the console prompt:

- 1 Go to the `installation_directory\inv\server\wminv\bin` directory.
- 2 At the prompt, enter `listser inventory_service_name`.

If the services are not up and running, check the Server Status log. For more information on the Server Status log, see [Section 18.3, “Viewing the Status of Inventory Components on an Inventory Server,” on page 712](#).

Stopping the Inventory Service on a Windows Inventory Server

To stop the Inventory services on a Windows 2000/2003 Inventory server:

- 1 In the Control Panel, double-click *Administrative Tools*.
- 2 Double-click *Services*.
- 3 Select *Novell Inventory Service*, then click *Stop*.

To stop a service on a Windows Inventory server from the console prompt:

- 1 Go to the `installation_directory\inv\server\wminv\bin` directory.
- 2 Enter `stopser inventory_service_name`.

13.1.5 Changing the Role of the Inventory Server

When you install ZENworks 7 Server Management, by default, the role of the Inventory server is a Standalone Server. By configuring the Inventory Service object, you can assign specific roles to the Inventory server based on your inventory deployment.

For example, if the deployment plan identifies three Inventory servers, such as a Root Server, an Intermediate Server with Database, and a Leaf Server for inventory deployment, you install Server Inventory on these servers, and choose the role for the Inventory server. Later, if you want to make changes in the inventory deployment, such as attaching the inventoried servers to the existing Root Server, you need to change the role of the Inventory Service object from Root Server to Root Server with Inventoried Servers. Additionally, depending on the new role, there are some policies you need to configure.

To change the role for any Inventory server:

- 1 Plan the change of roles carefully because the changes impact the existing inventory deployment. Also, consider the disk space requirements and ensure that you have the required configurations for Inventory.

- 2 In ConsoleOne, right-click the Inventory Service object (Inventory Service_ *server_name*), click *Properties*, then click the *Inventory Service Object Properties* tab.
- 3 Choose the new role of the Inventory Service object, then click *Apply*.
- 4 Bring down the services running on the changed Inventory server, follow the actions that you need to change the role, and then bring up the Inventory services.

To stop all Inventory Services:

- ♦ At NetWare server console prompt, enter the following commands:

```
stopser *
java -killZenWSInv
```
- ♦ On the Windows 2000/2003 server, from the Control Panel, double-click *Administrative Tools*, double-click *Services*, click *Novell Inventory Services*, then click *Stop*.

To restart all Inventory Services:

- ♦ At NetWare server console prompt, enter `startinv`
- ♦ On the Windows 2000/2003 server, from the Control Panel, double-click *Administrative Tools*, double-click *Services*, click *Novell Inventory Services*, then click *Start*.

The following sections contain information to help you change the role of the Inventory Service object:

- ♦ [“Changing the Role of the Root Server” on page 484](#)
- ♦ [“Changing the Role of the Root Server with Inventoried Servers” on page 485](#)
- ♦ [“Changing the Role of the Intermediate Server” on page 486](#)
- ♦ [“Changing the Role of the Intermediate Server with Database” on page 487](#)
- ♦ [“Changing the Role of the Intermediate Server with Database and Inventoried Servers” on page 488](#)
- ♦ [“Changing the Role of the Intermediate Server with Inventoried Servers” on page 489](#)
- ♦ [“Changing the Role of the Leaf Server” on page 490](#)
- ♦ [“Changing the Role of the Leaf Server with Database” on page 491](#)
- ♦ [“Changing the Role of the Standalone Server” on page 492](#)

Changing the Role of the Root Server

To change the role of the Root Server to a different role, perform the actions specified in [Table 13-3](#):

Table 13-3 Tasks to be performed to change the role of the Root Server

| To change the role of the Root Server to ... | Tasks: |
|--|---|
| Root Server with Inventoried Servers | Perform the following task: <ol style="list-style-type: none"> 1. After changing the role, configure the Server Inventory policy so that the inventoried servers that you have attached to the Root Server with Inventoried servers are scanned for. |

| To change the role of the Root Server to ... | Tasks: |
|--|---|
| Intermediate Server | Perform the following tasks: <ol style="list-style-type: none"> 1. Before changing the role, remove the Database Location policy associated with a Root Server. 2. After changing the role, configure the Roll-Up policy to specify the next-destination server for roll-up of information from this Inventory server. |
| Intermediate Server with Database | Perform the following task: <ol style="list-style-type: none"> 1. After changing the role, configure the Roll-Up policy to specify the next-destination Inventory server for roll-up of information from this Inventory server. |
| Intermediate Server with Database and Inventoried Servers | Perform the following tasks after changing the role: <ol style="list-style-type: none"> 1. Configure the Server Inventory policy so that the inventoried servers that you have attached are scanned for. 2. Configure the Roll-Up policy to specify the next-destination server for roll-up of information from this Inventory server. |
| Intermediate Server with Inventoried Servers | Perform the following tasks: <ol style="list-style-type: none"> 1. Before changing the role, remove the Database Location policy associated with the Root Server. 2. After changing the role, configure the Server Inventory policy so that the inventoried servers that you have attached are scanned for. 3. After changing the role, configure the Roll-Up policy to specify the next-destination Inventory server for roll-up of information from this Inventory server. |
| Leaf Server, Leaf Server with Database, or Standalone Server | Server Inventory does not allow you to change the Root Server to these Inventory servers because these changes affect the complete inventory system. If you want to assign these roles, you should reinstall and set up the Server Inventory component. |

Changing the Role of the Root Server with Inventoried Servers

Perform the actions specified in [Table 13-4](#):

Table 13-4 Tasks to be performed to change the role of the Root Server with Inventoried Servers

| To change the role of the Root Server with Inventoried Servers to ... | Tasks: |
|---|---|
| Root Server | Perform the following task: <ol style="list-style-type: none"> 1. Before changing the role, remove the Server Inventory policy associated with the Root Server with Inventoried Servers. |

| To change the role of the Root Server with Inventoried Servers to ... | Tasks: |
|---|--|
| Intermediate Server | Perform the following tasks: <ol style="list-style-type: none"> <li data-bbox="630 359 1411 415">1. Before changing this role, remove the Database Location policy and the Server Inventory policy. <li data-bbox="630 428 1414 485">2. After changing the role, configure the Roll-Up policy to specify the next-destination server for roll-up of information from this Inventory server. |
| Intermediate Server with Database | Perform the following tasks: <ol style="list-style-type: none"> <li data-bbox="630 564 1411 674">1. Before changing the role, if the Server Inventory policy is associated with the Root Server with Inventory servers, remove the policy for those servers attached to this Inventory server or to the lower-level Inventory servers that roll up to this Inventory server. <li data-bbox="630 686 1414 743">2. After changing the role, configure the Roll-Up policy to specify the next-destination server for roll-up of information from this Inventory server. |
| Intermediate Server with Database and Inventoried Servers | Perform the following task: <ol style="list-style-type: none"> <li data-bbox="630 823 1414 905">1. After changing the role, configure the Roll-Up policy to specify the next-destination Inventory server for roll-up of information from this Inventory server. |
| Intermediate Server with Inventoried Servers | Perform the following task: <ol style="list-style-type: none"> <li data-bbox="630 984 1403 1037">1. Before changing the role, remove the Database Location policy that is associated with the Root Server with Inventoried Servers. |
| Leaf Server, Leaf Server with Database, or Standalone Server | Server Inventory does not allow you to change the Root Server to these Inventory servers because these changes affect the complete inventory system. If you want to assign these roles, you should reinstall and set up the Server Inventory component. |

Changing the Role of the Intermediate Server

Perform the actions specified in [Table 13-5](#):

Table 13-5 Tasks to be performed to change the role of the Intermediate Server

| To change the role of the Intermediate Server to ... | Tasks: |
|--|---|
| Root Server | Perform the following tasks: <ol style="list-style-type: none"> <li data-bbox="630 1566 1214 1587">1. Before changing the role, remove the Roll-Up policy. <li data-bbox="630 1600 1333 1633">2. After changing the role, configure the Database Location policy. |
| Root Server with Inventory Servers | Perform the following tasks: <ol style="list-style-type: none"> <li data-bbox="630 1713 1214 1734">1. Before changing the role, remove the Roll-Up policy. <li data-bbox="630 1747 1403 1833">2. After changing the role, configure the Server Inventory policy for those inventoried servers attached to this server and the Database Location policy. |

| To change the role of the Intermediate Server to ... | Tasks: |
|--|--|
| Intermediate Server with Database | Perform the following task: <ol style="list-style-type: none"> 1. After changing the role, configure the Database Location policy for this Inventory server. |
| Intermediate Server with Database and Inventoried Servers | Perform the following tasks: <ol style="list-style-type: none"> 1. After changing the role, configure the Server Inventory policy so that all the inventoried servers associated to this Inventory Service object, and also those inventoried servers associated to the lower-level Inventory servers that roll up to this Inventory server are scanned for. 2. After changing the role, configure the Database Location policy. |
| Intermediate Server with Inventoried Servers | Perform the following task: <ol style="list-style-type: none"> 1. After changing the role, configure the Server Inventory policy so that the inventoried servers that you have attached are scanned for. |
| Leaf Server, Leaf Server with Database, or Standalone Server | Server Inventory does not allow you to change the Intermediate Server to these Inventory servers because these changes affect the complete inventory system. If you want to assign these roles, you should reinstall and set up the Server Inventory component. |

Changing the Role of the Intermediate Server with Database

Perform the actions specified in [Table 13-6](#):

Table 13-6 Tasks to be performed to change the role of the Intermediate Server with Database

| To change the role of the Intermediate Server with Database to ... | Tasks: |
|---|---|
| Root Server | Perform the following task: <ol style="list-style-type: none"> 1. Before changing the role, remove the Roll-Up policy associated with the Intermediate Server with Database. |
| Root Server with Inventoried Servers | Perform the following tasks: <ol style="list-style-type: none"> 1. Before changing the role, remove the Roll-Up policy associated with the Intermediate Server with Database. 2. After changing the role, configure the Server Inventory policy so that the inventoried servers that you have attached are scanned for. |
| Intermediate Server | Perform the following task: <ol style="list-style-type: none"> 1. Before changing the role, remove the Database Location policy that is associated with the Intermediate Server with Database. |
| Intermediate Server with Database and Inventoried Servers | Perform the following task: <ol style="list-style-type: none"> 1. After changing the role, configure the Server Inventory policy so that the inventoried servers attached are scanned for. |

To change the role of the Intermediate Server with Database to ...**Tasks:**

Intermediate Server with Inventoried Servers

Perform the following tasks:

1. Before changing the role, remove the Database Location policy that is associated with the Intermediate Server with Database.
2. After changing the role, configure the Server Inventory policy so that the inventoried servers that you have attached are scanned for.

Leaf Server, Leaf Server with Database, or Standalone Server

Server Inventory does not allow you to change the Intermediate Server to these Inventory servers because these changes affect the complete inventory system. If you want to assign these roles, you should reinstall and set up the Server Inventory component.

Changing the Role of the Intermediate Server with Database and Inventoried Servers

Perform the actions specified in [Table 13-7](#):

Table 13-7 Tasks to be performed to change the role of the Intermediate Server with Database and Inventoried Servers

To change the role of the Intermediate Server with Database and Inventoried Servers to ...**Tasks:**

Root Server

Perform the following tasks before changing the role:

1. Remove the Roll-Up policy associated with the Intermediate Server with Database and Inventoried Servers.
2. Remove the Server Inventory policy associated with the inventoried server so that the inventoried servers do not send the scan files to this server.

Root Server with Inventoried Servers

Perform the following task:

1. Before changing the role, remove the Roll-Up policy associated with the Intermediate Server with Database and Inventoried Servers.

Intermediate Server

Perform the following tasks before changing the role:

1. Remove the Server Inventory policy associated with the lower-level servers that roll up to the Intermediate Server with Database and Inventoried Servers.
2. Remove the Database Location policy associated with the Intermediate Server with Database and Inventoried Servers.

Intermediate Server with Database

Perform the following task:

1. Remove the Server Inventory policy of the Intermediate Server with Database and Inventoried Servers or reconfigure the policy.
-

To change the role of the Intermediate Server with Database and Inventoried Servers to ...

Tasks:

Intermediate Server with Inventoried Servers

Perform the following task:

1. Before changing the role, remove the Database Location policy associated with the Intermediate Server with Database and Inventoried Servers.

Leaf Server, Leaf Server with Database, Standalone Server

Server Inventory does not allow you to change the Intermediate Server to these servers because these changes affect the complete inventory system. If you want to assign these roles, you should reinstall and set up the Server Inventory component.

Changing the Role of the Intermediate Server with Inventoried Servers

Perform the actions specified in [Table 13-8](#):

Table 13-8 Tasks to be performed to change the role of the Intermediate Server with Inventoried Servers

To change the role of the Intermediate Server with Inventoried Servers to ...

Tasks:

Root Server

Perform the following tasks:

1. Before changing the role, remove the Roll-Up policy associated with the Intermediate Server with Inventoried Servers.
2. Before changing the role, remove the Server Inventory policy associated with the inventoried server so that the inventoried servers attached do not send the scan files to this Inventory server.
3. After changing the role, configure the Database Location policy for this Inventory server.

Root Server with Inventoried Servers

Perform the following tasks:

1. Before changing the role, remove the Roll-Up policy associated with the Intermediate Server with Inventoried Servers.
2. After changing the role, configure the Server Inventory policy for those inventoried servers attached to the lower-level Inventory server that roll up to this Inventory server.
3. After changing the role, configure the Database Location policy.

Intermediate Server

Perform the following task:

1. Before changing the role, remove the Server Inventory policy.

Intermediate Server with Database

Perform the following tasks:

1. Before changing the role, remove the Server Inventory policy associated to the inventoried server attached to this Inventory Service object.
 2. After changing the role, configure the Database Location policy for this Inventory server.
-

To change the role of the Intermediate Server with Inventoried Servers to ...**Tasks:**

Intermediate Server with Database and Inventoried Servers

Perform the following task:

1. After changing the role, configure the Database Location policy for this Inventory server.

Leaf Server, Leaf Server with Database or Standalone Server

Server Inventory does not allow you to change the Intermediate Server to these Inventory servers because these changes affect the complete inventory system. If you want to assign these roles, you should reinstall and set up the Server Inventory component.

Changing the Role of the Leaf Server

Perform the actions specified in [Table 13-9](#):

Table 13-9 Tasks to be performed to change the role of the Leaf Server

To change the role of the Leaf Server to ...**Tasks:**

Root Server

Perform the following tasks:

1. Before changing the role, remove the Roll-Up policy associated with the Leaf Servers.
2. Before changing the role, remove the Server Inventory policy associated with the inventoried server.
3. After changing the role, configure the Database Location policy for the Root Server.

Root Server with Inventoried Servers

Perform the following tasks:

1. Before changing the role, remove the Roll-Up policy associated with the Leaf Server.
2. After changing the role, configure the Database Location policy for the Root Server with Inventoried Servers.

Intermediate Server

Perform the following task:

1. Before changing the role, remove the Server Inventory policy for those inventoried servers associated with the Inventory server or reconfigure the policy.

Intermediate Server with Database

Perform the following tasks:

1. Before changing the role, remove the Server Inventory policy for those inventoried servers associated with the lower-level Inventory servers that roll up to this Inventory server or reconfigure the policy.
2. After changing the role, configure the Database Location policy for this Inventory server.

Intermediate Server with Database and Inventoried Servers

Perform the following task:

1. After changing the role, configure the Database Location policy for this Inventory server.
-

| To change the role of the Leaf Server to ... | Tasks: |
|---|---|
| Intermediate Server with Inventoried Servers | This change of role does not require any specific policy modifications. |
| Leaf Server with Database | Perform the following task: <ol style="list-style-type: none"> 1. After changing the role, configure the Database Location policy for this Inventory server. |
| Standalone Server | Perform the following task: <ol style="list-style-type: none"> 1. Before changing the role, remove the Roll-Up policy associated with the Leaf Server. |

Changing the Role of the Leaf Server with Database

Perform the actions specified in [Table 13-10](#):

Table 13-10 Tasks to be performed to change the role of the Leaf Server with Database

| To change the role of the Leaf Server with Database to ... | Tasks: |
|---|---|
| Root Server | Perform the following tasks before changing the role: <ol style="list-style-type: none"> 1. Remove the Server Inventory policy associated with the Leaf Server with Database. 2. Remove the Roll-Up policy associated with the Leaf Server with Database. |
| Root Server with Inventoried Servers | Perform the following task: <ol style="list-style-type: none"> 1. Before changing the role, remove the Roll-Up policy associated with the Leaf Server with Database. |
| Intermediate Server | Perform the following task: <ol style="list-style-type: none"> 1. Before changing the role, remove the Server Inventory policy and the Database Location policy associated with the Leaf Server with Database. |
| Intermediate Server with Database | Perform the following task: <ol style="list-style-type: none"> 1. Before changing the role, remove the Server Inventory policy associated with the Leaf Server with Database. |
| Intermediate Server with Database and Inventoried Servers | This change of role does not require any specific policy modifications. |
| Intermediate Server with Inventoried Servers | Perform the following task: <ol style="list-style-type: none"> 1. Before changing the role, remove the Database Location policy associated with the Leaf Server with Database. |

To change the role of the Leaf Server with Database to ...**Tasks:**

| | |
|-------------------|---|
| Leaf Server | Perform the following task: <ol style="list-style-type: none">1. Before changing the role, remove the Database Location policy associated with the Leaf Server with Database. |
| Standalone Server | Perform the following task: <ol style="list-style-type: none">1. Before changing the role, remove the Roll-Up policy associated with the Leaf Server with Database. |

Changing the Role of the Standalone Server

Perform the actions specified in [Table 13-11](#):

Table 13-11 Tasks to be performed to change the role of the Standalone Server

To change the role of the Standalone Server to ...**Tasks:**

| | |
|---|---|
| Root Server | Perform the following task: <ol style="list-style-type: none">1. Before changing the role, remove the Server Inventory policy associated with the Standalone Server. |
| Root Server with Inventoried Servers | This change of role does not require any specific policy modifications. |
| Intermediate Server | Perform the following tasks: <ol style="list-style-type: none">1. Before changing the role, remove the Server Inventory policy and the Database Location policy associated with the Standalone Server.2. After changing the role, configure the Roll-Up policy to specify the next-destination Inventory server for roll-up of information from the Intermediate Server with Database. |
| Intermediate Server with Database | Perform the following tasks: <ol style="list-style-type: none">1. Before changing the role, remove the Server Inventory policy associated with the Standalone Server.2. After changing the role, configure the Roll-Up policy to specify the next-destination Inventory server for roll-up of information from the Intermediate Server with Database. |
| Intermediate Server with Database and Inventoried Servers | Perform the following task: <ol style="list-style-type: none">1. After changing the role, configure the Roll-Up policy to specify the next-destination server for roll-up of information from the Intermediate Server with Database and Inventoried Servers. |

To change the role of the Standalone Server to ...

Tasks:

Intermediate Server with Inventoried Servers

Perform the following tasks:

1. Before changing the role, remove the Database Location policy associated with the Standalone Server.
2. After changing the role, configure the Roll-Up policy to specify the next-destination server for roll-up of information from the Intermediate Server with Inventoried Servers.

Leaf Server

Perform the following tasks:

1. Before changing the role, remove the Database Location policy associated with the Standalone Server.
2. After changing the role, configure the Roll-Up policy to specify the next-destination server for roll-up of information from the Leaf Server.

Leaf Server with Database

Perform the following task:

1. After changing the role, configure the Roll-Up policy to specify the next-destination server for roll-up of information from the Leaf Server with Database.
-

13.2 Setting Up the Inventory Database

The following sections contain detailed information to help you set up your Inventory database:

- ♦ [Section 13.2.1, “Setting Up the Sybase Inventory Database,” on page 493](#)
- ♦ [Section 13.2.2, “Setting Up the Oracle Inventory Database,” on page 500](#)
- ♦ [Section 13.2.3, “Setting Up the MS SQL Server 2000 or MS SQL Server 2005 Inventory Database,” on page 510](#)

If you want to replace the Inventory database, always stop the Inventory services before replacing the database. Replace the database and restart the Inventory services. For more information, see [Section 13.1.4, “Starting and Stopping the Inventory Service,” on page 482](#).

13.2.1 Setting Up the Sybase Inventory Database

This section contains the following information:

- ♦ [“Adding Non-English Enumerated Values for Inventory Attributes into the Inventory Database” on page 494](#)
- ♦ [“Manually Creating the Sybase Inventory Database Object” on page 494](#)
- ♦ [“Organizing the Sybase Database Spaces on NetWare or Windows Servers \(AlterDBSpace Tool\)” on page 495](#)
- ♦ [“Understanding the Sybase Database Startup Parameters” on page 497](#)
- ♦ [“Backing Up the Sybase Inventory Database” on page 498](#)

Adding Non-English Enumerated Values for Inventory Attributes into the Inventory Database

You need to add the non-English enumerated (enum) values so the Inventory report can display the enum value for the inventory attributes in internationalized versions. The non-English enum values must be available in English version of the product so that the rolled-up inventory information from non-English sites can be properly captured at the high-level servers where only English versions are installed.

For more information about the list of attributes that contains enumerated values, see [Appendix M, "Enumeration Values,"](#) on page 767.

To add the non-English enum values:

- 1 Specify the JDBC connection settings in the `connection.prop` to connect to the Sybase database. The file is located in the following NetWare or Windows directory:

```
zenworks_directory\inv\server\wminv\ properties
```

You can do this by copying the template property settings for Sybase specified in the comments section in the `connection.prop` file. Specify the IP address, port number, and Database SID in the JDBC URL string that matches your Sybase server configuration.

- 2 At the server prompt, enter `AddEnums`
`directory_name_containing_connection.prop`

On the NetWare or Windows Inventory server, run the above command from the following directory:

```
zenworks_directory\inv\server\wminv\bin
```

After executing the command, the a message indicating that the non-English enums have been successfully inserted is displayed on the console prompt.

Manually Creating the Sybase Inventory Database Object

To manually create the Inventory Database object (Inventory database_*server_name*) for Sybase:

- 1 In ConsoleOne, right-click in the eDirectory tree where you want to create the database object, click *New*, click *Object*, click *ZENworks Database*, then click *OK*.
- 2 Enter a name for the database object, then click *OK*.
- 3 Configure the Database server options of the Database object.
 - 3a In ConsoleOne, right-click the database object, click *Properties*, then click the *ZENworks Database* tab.
 - 3b Select the database server object using any of the following methods:

- ♦ If eDirectory is installed on the database server, in the *Server DN* field, browse for and select the object for the server where the database is physically installed and running.

The server's IP address is automatically populated to the *Server IP Address or DNS Name* drop-down list. If the selected server object has more than one IP address, select the appropriate IP address.

To clear the value set in the *Server DN* field, type the IP address of another database server or browse and select another server object.

- ♦ If eDirectory is not installed on the database server, then enter the server's IP address or the DNS name in the *Server IP Address or DNS Name* field.

3c Type the values for the following options:

- ♦ **Database (Read-Write) Username:** *MW_DBA*
- ♦ **Database (Read-Write) Password:** *novell*
- ♦ **Database (Read Only) Username:** *MW_READER*
- ♦ **Database (Read Only) Password:** *novell*
- ♦ **Database (Write Only) Username:** *MW_UPDATER*
- ♦ **Database (Write Only) Password:** *novell*

IMPORTANT: All Inventory components use the username and password the configured in the database object. By default, “novell” is the password for all options. But you can change it in the database, and update the same here.

3d Click *Apply*.

3e To configure the JDBC* Driver properties, click the *Jdbc Driver Information* tab.

3f Select *Sybase*, then click *Default Settings*.

This populates the fields with default JDBC driver information.

The database settings for Sybase are:

- ♦ **Driver:** *com.sybase.jdbc.SybDriver*
- ♦ **Protocol:** *jdbc:*
- ♦ **SubProtocol:** *sybase:*
- ♦ **SubName:** *Tds:*
- ♦ **Port:** *2638*
- ♦ **Flags:** *?ServiceName=mgmtdb&JCONNECT_VERSION=4*
- ♦ **Database Service Name:** *the database name specified against the -n Sybase startup parameter while invoking Sybase.*

NOTE: By default, the value of the *-n* switch is the IP address of the database server. If you retain this switch value, you must enter the same IP address as the database service name.

3g Click *Apply*, then click *Close*.

Organizing the Sybase Database Spaces on NetWare or Windows Servers (AlterDBSpace Tool)

If there are more volumes or drives on the multiple physical disks of the database server, placing the Sybase database spaces files on separate volumes or drives improves the performance while accessing the database.

If you install the Sybase database component of ZENworks 7 Server Management, the system database file and the database spaces files are installed in the location on the database server you specify. On loading the Inventory database server, the system database file (*mgmtdb.db*) is loaded. This *mgmtdb.db* file references the inventory information in the database spaces files. The database spaces files (*mgmtdb1.db*, *mgmtdb2.db*, *mgmtdb3.db*, *mgmtdb4.db*,

mgmtdb5.db, mgmtdb6.db, mgmtdb7.db, mgmtdb8.db, mgmtdb9.db, mgmtdb10.db, and mgmtdb11.db) contain the inventory information.

The `alterdb.props` file is installed on the database server in the following NetWare or Windows location:

```
inventory_server_installation_directory\wminv\properties
```

You can modify the sections in the file to specify the location of the database spaces on the volumes or drives.

The contents of the `alterdb.props` file are as follows:

```
#Database Space Properties
count=11
mgmtdb1=location_of_mgmtdb1
mgmtdb2=location_of_mgmtdb2
mgmtdb3=location_of_mgmtdb3
mgmtdb4=location_of_mgmtdb4
mgmtdb5=location_of_mgmtdb5
mgmtdb6=location_of_mgmtdb6
mgmtdb7=location_of_mgmtdb7
mgmtdb8=location_of_mgmtdb8
mgmtdb9=location_of_mgmtdb9
mgmtdb10=location_of_mgmtdb10
mgmtdb11=location_of_mgmtdb11
```

.....

To organize the database spaces:

- 1 Ensure that the database is not loaded.
- 2 Ensure that the Inventory Service Manager is not running on the Inventory server.
- 3 Manually move the database spaces files on the Inventory server.

Arrange the database spaces files as follows for better performance:

- ♦ MGMTDB1 and MGMTDB2 in the same location
- ♦ MGMTDB3 and MGMTDB6 in the same location
- ♦ MGMTDB5 and MGMTDB7 in the same location
- ♦ MGMTDB8 and MGMTDB4 in the same location
- ♦ MGMTDB9 and MGMTDB10 in the same location
- ♦ MGMTDB11 in a location

IMPORTANT: If you move `mgmtdb.db` to another directory or volume on a NetWare server, update the `sys:\system\mgmt dbs.ncf` file with the new location of the `mgmtdb.db`.

If you move `mgmtdb.db` to another directory or volume on a Windows server, run the `ntdbconfig.exe` located in the `zenworks\database\dbengine` directory. In the NTDBCONFIG dialog box, enter the new path of `mgmtdb.db`.

- 4 Modify the location of the eleven database space files in the `alterdb.props` file.

For example,

- ♦ On NetWare: Enter `mgmt.db3=SYS: \\ZENWORKS \\ INV \\ DB`
- ♦ On Windows: Enter `mgmt.db3=C: \\ZENWORKS \\ INV \\ DB`

5 Load the database.

- ♦ On NetWare: Enter `mgmt.dbs`.
- ♦ On Windows: Run the database service.

Ignore the error messages displayed on the console. These messages are displayed because the database spaces files are not loaded.

6 Ensure that the Database Location policy has been configured.

7 On the NetWare or Windows Inventory server console, run the AlterDBSpace service by entering the following command at the server prompt:

```
StartSer AlterDBSpace
```

On the Inventory server, the AlterDBSpace tool runs as a service.

You will see a message that the database is adjusted.

8 Exit the database and then load the database.

Ensure that there are no errors while loading the database. Errors indicate that the specified location of the database spaces files is incorrect or does not exist. Ensure that the path to the database spaces files is correct in the `alterdb.props` file and repeat the procedure to organize the database spaces files.

IMPORTANT: If you place the database spaces files in different volumes or drives, the log file should be placed in the same volume or drive as the system database file (`mgmt.db.db`).

Understanding the Sybase Database Startup Parameters

The startup parameters of the Sybase database are as follows:

- ♦ **-c:** Sets the initial memory reserves for caching database pages and other server information. For example, `-c 32M` reserves 32 MB cache size.
- ♦ **-gc:** Sets the maximum length of time in minutes that the database server runs without doing a checkpoint on each database. The default value is 60 minutes. For example, `-gc` sets the checkpoint time as 120 minutes.
- ♦ **-gn:** Sets the number of execution threads to be used in the database server.
- ♦ **-m:** Deletes the transaction log when a checkpoint is done, either at shutdown or as a result of a checkpoint scheduled by the server.
- ♦ **-n:** Specifies the host name of the database server. For example, `-n IP_address`.
- ♦ **-ti:** Disconnects the connections that have not submitted a request for a certain number of minutes. The default is 240 (4 hours). A client machine in the middle of the database transaction locks until the transaction ends or the connection terminates. The `-ti` option is provided to disconnect inactive connections and to free their locks. For example, specify `-ti 400`.
- ♦ **-x:** Specifies a communication link. For example, `-x tcpip` indicates a TCP/IP link.
- ♦ **-ct:** Enables character set translation by converting strings between character sets that represent the same characters but at different values. This is useful when the client machine and the database use different character sets.

- ◆ **-gss:** Sets the stack size per internal execution thread in the server.
- ◆ **database_installation_path:** Specifies the installation path of the Inventory database. For example, `c:\zenworks\inv\db\mgmtdb.db`.

Backing Up the Sybase Inventory Database

Server Inventory provides a utility, Database Backup, to back up the Sybase Inventory database from the server. We recommend that you back up the database on a weekly basis. However, if you are tracking the inventory of servers frequently, increase the frequency of backup.

You can back up the database files and the transaction log to the location relative to the `\scandir` path.

You can run Database Backup either from the server console or ConsoleOne.

This section provides information on the following topics:

Running Database Backup from the Server Console

Before running Database Backup from the server console, fulfill the following prerequisites:

- ❑ You can run Database Backup only on an Inventory server to which you have associated a database server. If you deployed more than one database server, you must run Database Backup for each database server.
- ❑ Ensure that the database you have to back up is configured in the Database Location policy. For more information on how to access the Database Location policy, see [Section 13.4, “Configuring the Database Location Policy,” on page 521](#).
- ❑ The backup files are relative to the SCANDIR path. For example, if the SCANDIR path is `sys:\zenworks\inv\scandir`, the database will be backed up in `sys:\zenworks\inv\scandir\Backup` directory by default. To backup in another directory, in the [DBBackup Service] section of the server property file, you must modify the location of the backup destination in the ARGUMENTS parameter. For example, if the value of ARGUMENTS parameter is changed to `“Backup\day1”`, the database is backed up in `sys:\zenworks\inv\scandir\backup\day1`.

NOTE: If the directory path is multi-level, enclose the value in two double quotes and use `\\` instead of `\` as path separator.

You must modify the server property file located on the server on which you are running Database Backup. Modify the server property file corresponding to the role of the server. For example, if you are running Database Backup on the Leaf Server with Database, modify the server property file, `leaf_db_wks.properties`.

IMPORTANT: When the properties file is modified, stop and start the inventory service for the modified property file to be loaded.

- ❑ Ensure that the Service Manager is loaded when you run Database Backup.

To run Database Backup from the server console:

- 1 At the Inventory server console, enter `StartSer DBBACKUP`.

- 2 View the status of the backup in the backup log file. The database is copied to `zenworks_installation_path\zenworks\inv\scandir\directory_you_specify`.

Database Backup creates a log file, `bacstatus.txt`, located in the `zenworks\inv\scandir` directory on NetWare and Windows 2000 servers. The log records the status of the backup operation. Open this text file to view the status of the backup. This file increases in size for every backup operation. Remove the existing contents of the file if you do not require the details.

Running Database Backup from ConsoleOne

- 1 In ConsoleOne, click *Tools*, click *ZENworks Inventory*, then click *Database Backup*.

If you want to back up the latest information in the Inventory database, right-click the database object, click *ZENworks Inventory*, then click *Database Backup*.

- 2 Enter the path to the directory where the database backup will be saved

WARNING: Do not use double-byte characters in the directory name. If you do so, Sybase interprets the double-byte characters as a different name, and backs up the database in the directory with the interpreted name.

If the Inventory database is running on a NetWare server, you can either enter the path or click *Browse* to browse for and select a directory. If you enter the database backup directory name without specifying the complete path, the backup directory will be created in the `sys :` directory.

If the Inventory database is running on a Windows machine, you must manually enter the backup directory path. If you enter the database backup directory name without specifying the complete path, the backup directory is created in the `\winnt\system32` directory on Windows, and in the `root` directory on Linux.

NOTE: If you want to back up the database to a non-existent directory, only one level of the new directory will be created. To back up the database to a subdirectory, ensure that the primary directory already exists. For example, if you want to back up the database to a new `c : \backup` directory, the `\backup` directory will be created and the database will be backed up. But if you want to back up the database to a new `\database` directory, located under `c : \backup`, the `\backup` directory must already exist.

- 3 Click *Start Backup*.

This backs up the database to the specified directory on the server running the database and overwrites any existing files without prompting about the overwrite.

Database Backup creates a log file, `Backupst.txt`, located in the `ConsoleOne_installation_directory\1.2\bin` directory. The log records the status of the backup operation. Open this text file to view the status of the backup. This file increases in size for every backup operation. Remove the existing contents of the file if you do not require the details.

Restoring the Inventory Database

- 1 If the Inventory database server is up, stop the Storer service. At the database server console, enter `StopSer Storer`.

- 2 Exit the Sybase database.
 - ♦ On NetWare servers: At the database server prompt, enter `q` to stop the Sybase database.
 - ♦ On Windows 2000/2003: In the Windows Control Panel, double-click *Administrative Tools*, double-click *Services*, select *Novell Database - Sybase*, then click *Stop*.
- 3 Copy the backup files, overwriting the working database files.
- 4 Restart the database server.

13.2.2 Setting Up the Oracle Inventory Database

The following sections explain how to set up the Inventory database for Oracle9i and Oracle10g:

- ♦ [“Creating the Oracle9i Inventory Database on a Windows Server” on page 500](#)
- ♦ [“Creating the Oracle9i Inventory Database on a UNIX Server” on page 502](#)
- ♦ [“Creating the Oracle10g Inventory Database on a Windows Server” on page 504](#)
- ♦ [“Creating the Oracle10g Inventory Database on a UNIX Server” on page 506](#)
- ♦ [“Manually Creating the Oracle Inventory Database Object” on page 508](#)
- ♦ [“Configuring and Running Multiple Oracle Database Instances on a Windows Server” on page 509](#)

IMPORTANT: In this setup, the Inventory database is not mounted with any other version or instances of Oracle databases.

Creating the Oracle9i Inventory Database on a Windows Server

Make sure that the following prerequisites are met:

- Oracle 9.2.0.6 must be installed on the server before configuring the Inventory database.
- To maintain the Inventory database on Oracle, Server Inventory requires that you have a minimum of 25 user licenses.

You must manually create the Inventory database for Oracle on Windows servers by following the procedure below:

- 1 Create a directory `c:\schema` and copy all the files in the following directories from the *ZENworks 7 Companion 2 CD* to the `schema` directory:
 - `database\oracle9i\common`
 - `database\oracle9i\winntspecific`
- 2 By default, `_create.sql`, `init.ora`, `_start.sql` are Read-only files. Make the files writable.
- 3 Create the `user_specified_path\zenworks\inventory\oracle\database\trace` directory structure.
- 4 In `c:\schema_create.sql`, replace all instances of `d:` with `user_specified_path`.
- 5 In `c:\schema\init.ora`, replace all instances of `d:` with `user_specified_path`.
- 6 In `c:\schema_start.sql`, replace all instances of `d:` with `user_specified_path`.
If `d:` is not found, check and correct the path of `init.ora` in the database directory.

- 7 Copy `c:\schema\init.ora` to `user_specified_path\zenworks\inventory\oracle\database`.
- 8 Copy `c:\schema_start.sql` to `user_specified_path\zenworks`.
- 9 Make sure that Oracle services are loaded correctly and the database is not mounted.
- 10 At the command prompt, enter `sqlplus /nolog` to load the Oracle server manager.
- 11 At the Oracle Server Manager prompt (sqlplus prompt), enter `@c:\schema\schema.sql`.
Review the `c:\schema\inv.log` file to make sure that the database has been created successfully. If the database has not been successfully created, `inv.log` contains the following error messages: Oracle not available, Out of space, Compilation error.
- 12 Add non-English enumerated (enum) values for certain Inventory attributes into the Inventory database.

IMPORTANT: You must perform this step on the English version of the product also.

You need to add the non-English enumerated values so the Inventory ConsoleOne utilities such as Inventory Report can display the enum value for the inventory attributes in internationalized versions. The non-English enum values must be available in English version of the product so that the rolled-up inventory information from non-English sites can be properly captured at the high-level servers where only English versions are installed.

For more information about the list of attributes that contains enumerated values, see [Appendix M, “Enumeration Values,” on page 767](#).

To add the non-English enum values:

- 12a Specify the JDBC connection settings in the `zenworks_directory\inv\server\wminv\properties\connection.prop` file to connect to the Oracle database.
You can do this by copying the template property settings for Oracle specified in the comments section in the `connection.prop` file. Specify the IP address, port number, and Database SID in the JDBC URL string that matches your Oracle server configuration.
- 12b At the server prompt, enter `AddEnums directory_name_containing_connection.prop`
If your Inventory server is running on a Windows machine, run the above command from `zenworks_directory\inv\server\wminv\bin`.
- 12c Execute the following SQL statement at the sqlplus prompt to make sure that the localized enumerated values are added correctly:
 - ◆ To display the enumerated values in French: `connect mw_dba/password` and `SELECT * FROM cim.ostype_fr`
 - ◆ To display the enumerated values in Spanish: `connect mw_dba/password` and `SELECT * FROM cim.ostype_es;`
 - ◆ To display the enumerated values in German: `connect mw_dba/password` and `SELECT * FROM cim.ostype_de;`
 - ◆ To display the enumerated values in Brazilian-Portuguese: `connect mw_dba/password` and `SELECT * FROM cim.ostype_pt_BR;`
- 13 At the sqlplus prompt, enter `@c:\schema\schema1.sql`.
Review the `c:\schema\inv1.log` file to make sure that the database has been created successfully. If the database has not been successfully created, `inv1.log` contains the

following error messages: Oracle not available, Out of space, Compilation error.

- 14 At the sqlplus prompt, enter `connect / as sysdba` to login as DBA.
- 15 At the sqlplus prompt, enter `shutdown immediate`.
- 16 At the sqlplus prompt, enter `@path\zenworks_start.sql` to start the Inventory database.
- 17 Continue with [“Manually Creating the Oracle Inventory Database Object” on page 508](#).

Creating the Oracle9i Inventory Database on a UNIX Server

Make sure that the following prerequisites are met:

- Oracle 9.2.0.6 must be installed must be installed on Linux or Solaris versions supported by Oracle9i.
- Hard disk free space: 4 GB or above.
- Primary memory: 1 GB or above.
- To maintain the Inventory database on Oracle, Server Inventory requires that you have a minimum of 25 user licenses.

You must manually create the Inventory database for Oracle on UNIX servers by following the procedure below:

- 1 Log in as an Oracle user.
- 2 Create a `/schema` directory in the Oracle installation directory (by default, Oracle is installed in the `/opt/oracle` directory), and copy all the files in the following directories from the *ZENworks 7 Companion 2 CD* to the schema directory:

```
database\oracle9i\common
database\oracle9i\unixspecific
```
- 3 By default, `_create.sql`, `init.ora`, `_start.sql` are Read-only files. Make the files writable.
- 4 Create the `user_specified_directory_path/zenworks/inventory/oracle/database/trace` directory structure in `/opt/oracle`.
- 5 In `schema/init.ora`, replace all instances of `$HOME` with `user_specified_directory_path`.
- 6 In `schema/_start.sql`, replace all instances of `$HOME` with `user_specified_directory_path`.
- 7 In `schema/_create.sql`, replace all instances of `$HOME` with `user_specified_directory_path`.
- 8 In `schema/schema.sql`, replace all instances of `$HOME` with the schema directory created in [Step 2](#).
- 9 In `schema/schema1.sql`, replace all instances of `$HOME` with the schema directory created in [Step 2](#).
- 10 Copy `schema/init.ora` to `user_specified_directory_path/zenworks/inventory/oracle/database`.
- 11 Copy `schema/_start.sql` to `user_specified_directory_path`.
- 12 Make sure the Oracle services are up and running and no database is mounted.

- 13 At the command prompt, enter `sqlplus /nolog` to load the Oracle Server Manager.
- 14 At the Oracle Server Manager prompt, enter `@$HOME/schema/schema.sql`, where `$HOME` is the schema directory created in [Step 2](#).
- 15 Review the `schema/inv.log` file to make sure that the database has been created successfully. If the database has not been successfully created, `inv.log` will contain the following error messages: Oracle not available, Out of space, Compilation error.
- 16 Add non-English enumerated (enum) values for certain Inventory attributes into the Inventory database.

IMPORTANT: You must perform this step on the English version of the product also.

You need to add the non-English enumerated values so the Inventory ConsoleOne utilities such as Inventory Report can display the enum value for the inventory attributes in internationalized versions. The non-English enum values must be available in English version of the product so that the rolled-up inventory information from non-English sites can be properly captured at the high-level servers where only English versions are installed.

To add the non-English enum values:

- 16a On the Inventory Server machine, Specify the JDBC connection settings in the `zenworks_directory\inv\server\wminv\properties\connection.prop` file to connect to the Oracle database.

You can do this by copying the template property settings for Oracle specified in the comments section in the `connection.prop` file. Specify the IP address, port number, and Database SID in the JDBC URL string that matches your Oracle server configuration.

- 16b At the server prompt, enter `AddEnums directory_name_containing_connection.prop`.

If your Inventory server is running on a Windows machine, run the above command from `zenworks_directory\inv\server\wminv\bin`.

- 16c Execute the following SQL statement at the `sqlplus` prompt to make sure that the localized enumerated values are added correctly:

- ♦ To display the enumerated values in French: `connect mw_dba/password` and `SELECT * FROM cim.ostype_fr`
- ♦ To display the enumerated values in Spanish: `connect mw_dba/password` and `SELECT * FROM cim.ostype_es;`
- ♦ To display the enumerated values in German: `connect mw_dba/password` and `SELECT * FROM cim.ostype_de;`
- ♦ To display the enumerated values in Brazilian-Portuguese: `connect mw_dba/password` and `SELECT * FROM cim.ostype_pt_BR;`
- ♦ To display the enumerated values in Japanese: `connect mw_dba/password` and `SELECT * FROM cim.ostype_ja;`

- 17 At the `sqlplus` prompt, enter `@$HOME/schema/schema1.sql`, where `$HOME` is schema directory created in [Step 2](#).

Review the `schema/inv1.log` file to make sure that the database has been created successfully. If the database has not been successfully created, `inv1.log` contains the following error messages: Oracle not available, Out of space, Compilation error.

- 18** At the sqlplus prompt, enter `connect / as sysdba` to login as DBA.
- 19** At the sqlplus prompt, enter `shutdown immediate`.
- 20** At the Oracle Server Manager prompt, enter `@user_specified_directory_path/zenworks/_start.sql` to start the Inventory database.

Creating the Oracle10g Inventory Database on a Windows Server

NOTE: Information about Oracle 10g R2 is applicable only for ZENworks 7 with Support Pack 1.

Make sure that the following prerequisites are met:

- Oracle10g R1 or Oracle10g R2 must be installed on the server before configuring the Inventory database.
- To maintain the Inventory database on Oracle, Server Inventory requires that you have a minimum of 25 user licenses.

You must manually create the Inventory database for Oracle on Windows servers by following the procedure below:

- 1** Create a directory `c:\schema`.
- 2** (Conditional) To create the Oracle10g R1 Inventory database, copy all the files in the following directories from the *ZENworks 7 Companion 2 CD* to the `c:\schema` directory:
 - `database\oracle10gR1\common`
 - `database\oracle10gR1\winntspecific`
- 3** (Conditional) To create the Oracle10g R2 Inventory database, copy all the files in the following directories from the *ZENworks 7 Companion 2 CD* to the `c:\schema` directory:
 - `database\oracle10gR2\common`
 - `database\oracle10gR2\winntspecific`
- 4** By default, `_create.sql`, `init.ora`, `_start.sql` are Read-only files. Make the files writable.
- 5** Create the `user_specified_path\zenworks\inventory\oracle\database\trace` directory structure.
- 6** In `c:\schema_create.sql`, replace all instances of `d:` with `user_specified_path`.
- 7** In `c:\schema\init.ora`, replace all instances of `d:` with `user_specified_path`.
- 8** In `c:\schema_start.sql`, replace all instances of `d:` with `user_specified_path`.
If `d:` is not found, check and correct the path of `init.ora` in the database directory.
- 9** Copy `c:\schema\init.ora` to `user_specified_path\zenworks\inventory\oracle\database`.
- 10** Copy `c:\schema_start.sql` to `user_specified_path\zenworks`.
- 11** Make sure that Oracle services are loaded correctly and the database is not mounted.
- 12** At the command prompt, enter `sqlplus /nolog` to load the Oracle server manager.
- 13** At the Oracle Server Manager prompt (sqlplus prompt), enter `@c:\schema\schema.sql`.
Review the `c:\schema\inv.log` file to make sure that the database has been created successfully. If the database has not been successfully created, `inv.log` contains the

following error messages: Oracle not available, Out of space, Compilation error.

- 14 Add non-English enumerated (enum) values for certain Inventory attributes into the Inventory database.

IMPORTANT: You must perform this step on the English version of the product also.

You need to add the non-English enumerated values so the Inventory ConsoleOne utilities such as Inventory Report can display the enum value for the inventory attributes in internationalized versions. The non-English enum values must be available in English version of the product so that the rolled-up inventory information from non-English sites can be properly captured at the high-level servers where only English versions are installed.

For more information about the list of attributes that contains enumerated values, see [Appendix M, “Enumeration Values,” on page 767](#).

To add the non-English enum values:

- 14a** On the Inventory Server machine, Specify the JDBC connection settings in the `connection.prop` file to connect to the Oracle database. The file is located in `zenworks_directory\inv\server\wminv\properties\` on Windows and in `/etc/opt/novell/zenworks/inv` on Linux.

You can do this by copying the template property settings for Oracle specified in the comments section in the `connection.prop` file. Specify the IP address, port number, and Database SID in the JDBC URL string that matches your Oracle server configuration.

- 14b** If your inventory server is running on a Windows machine, at the server prompt change to `zenworks_directory\inv\server\wminv\bin` and enter `AddEnums zenworks_directory\inv\server\wminv\properties`.

If your Inventory server is running on a Linux machine, at the server prompt change to `/opt/novell/bin/` and enter `AddEnums /etc/opt/novell/zenworks/inv`.

- 14c** Execute the following SQL statement at the sqlplus prompt to make sure that the localized enumerated values are added correctly:

- ♦ To display the enumerated values in French: `connect mw_dba/password` and `SELECT * FROM cim.ostype_fr`
- ♦ To display the enumerated values in Spanish: `connect mw_dba/password` and `SELECT * FROM cim.ostype_es;`
- ♦ To display the enumerated values in German: `connect mw_dba/password` and `SELECT * FROM cim.ostype_de;`
- ♦ To display the enumerated values in Brazilian-Portuguese: `connect mw_dba/password` and `SELECT * FROM cim.ostype_pt_BR;`

- 15 At the sqlplus prompt, enter `@c:\schema\schema1.sql`.

Review the `c:\schema\inv1.log` file to make sure that the database has been created successfully. If the database has not been successfully created, `inv1.log` contains the following error messages: Oracle not available, Out of space, Compilation error.

- 16 At the sqlplus prompt, enter `connect / as sysdba` to login as DBA.

- 17 At the sqlplus prompt, enter `shutdown immediate;`

- 18 At the sqlplus prompt, enter `@path\zenworks_start.sql` to start the Inventory database.

19 Continue with “[Manually Creating the Oracle Inventory Database Object](#)” on page 508.

Creating the Oracle10g Inventory Database on a UNIX Server

NOTE: Information about Oracle 10g R2 is applicable only for ZENworks 7 with Support Pack 1

Make sure that the following prerequisites are met:

- Oracle10g R1 or Oracle10gR2 must be installed must be installed on Linux or Solaris versions supported by Oracle10g.
- Hard disk free space: 4 GB or above.
- Primary memory: 1 GB or above.
- To maintain the Inventory database on Oracle, Server Inventory requires that you have a minimum of 25 user licenses.

You must manually create the Inventory database for Oracle on UNIX servers by following the procedure below:

- 1 Log in as an Oracle user.
- 2 Create a `/schema` directory in the Oracle installation directory (by default, Oracle is installed in the `/opt/oracle` directory).
- 3 (Conditional) To create the Oracle10g R1 Inventory database, copy all the files in the following directories from the *ZENworks 7 Companion 2 CD* to the `schema` directory:

```
Database\Oracle10gR1\Common
Database\Oracle10gR1\unixSpecific
```

- 4 (Conditional) To create the Oracle10g R2 Inventory database, copy all the files in the following directories from the *ZENworks 7 Companion 2 CD* to the `schema` directory:

```
Database\Oracle10gR2\Common
Database\Oracle10gR2\unixSpecific
```

- 5 By default, `_create.sql`, `init.ora`, `_start.sql` are Read-only files. Make the files writable.
- 6 Create the `user_specified_directory_path/zenworks/inventory/oracle/database/trace` directory structure in `/opt/oracle`.
- 7 In `schema/init.ora`, replace all instances of `$HOME` with `user_specified_directory_path`.
- 8 In `schema/_start.sql`, replace all instances of `$HOME` with `user_specified_directory_path`.
- 9 In `schema/_create.sql`, replace all instances of `$HOME` with `user_specified_directory_path`.
- 10 In `schema/schema.sql`, replace all instances of `$HOME` with the `schema` directory created in [Step 2](#).
- 11 In `schema/schema1.sql`, replace all instances of `$HOME` with the `schema` directory created in [Step 2](#).
- 12 Copy `schema/init.ora` to `user_specified_directory_path/zenworks/inventory/oracle/database`.
- 13 Copy `schema/_start.sql` to `user_specified_directory_path`.

- 14 Make sure the Oracle services are up and running and no database is mounted.
- 15 At the command prompt, enter `sqlplus /nolog` to load the Oracle Server Manager.
- 16 At the Oracle Server Manager prompt, enter `@$HOME/schema/schema.sql`, where `$HOME` is the schema directory created in [Step 2](#).
- 17 Review the `schema/inv.log` file to make sure that the database has been created successfully. If the database has not been successfully created, `inv.log` contains the following error messages: `Oracle not available`, `Out of space`, `Compilation error`.
- 18 Add non-English enumerated (enum) values for certain Inventory attributes into the Inventory database.

IMPORTANT: You must perform this step on the English version of the product also.

You need to add the non-English enumerated values so the Inventory ConsoleOne utilities such as Inventory Report can display the enum value for the inventory attributes in internationalized versions. The non-English enum values must be available in English version of the product so that the rolled-up inventory information from non-English sites can be properly captured at the high-level servers where only English versions are installed.

To add the non-English enum values:

- 18a** On the Inventory Server machine, Specify the JDBC connection settings in the `connection.prop` file to connect to the Oracle database. The file is located in `zenworks_directory\inv\server\wminv\properties\on` Windows and in `/etc/opt/novell/zenworks/inv` on Linux.

You can do this by copying the template property settings for Oracle specified in the comments section in the `connection.prop` file. Specify the IP address, port number, and Database SID in the JDBC URL string that matches your Oracle server configuration.

- 18b** If your inventory server is running on a Windows machine, at the server prompt change to `zenworks_directory\inv\server\wminv\bin` and enter `AddEnums zenworks_directory\inv\server\wminv\properties`.

If your Inventory server is running on a Linux machine, at the server prompt change to `/opt/novell/bin/` and enter `AddEnums /etc/opt/novell/zenworks/inv`.

- 18c** Execute the following SQL statement at the `sqlplus` prompt to make sure that the localized enumerated values are added correctly:

- ◆ To display the enumerated values in French: `connect mw_dba/password` and `SELECT * FROM cim.ostype_fr`
- ◆ To display the enumerated values in Spanish: `connect mw_dba/password` and `SELECT * FROM cim.ostype_es;`
- ◆ To display the enumerated values in German: `connect mw_dba/password` and `SELECT * FROM cim.ostype_de;`
- ◆ To display the enumerated values in Brazilian-Portuguese: `connect mw_dba/password` and `SELECT * FROM cim.ostype_pt_BR;`
- ◆ To display the enumerated values in Japanese: `connect mw_dba/password` and `SELECT * FROM cim.ostype_ja;`

- 19 At the `sqlplus` prompt, enter `@$HOME/schema/schema1.sql`, where `$HOME` is schema directory created in [Step 2](#).

Review the `schema/inv1.log` file to make sure that the database has been created successfully. If the database has not been successfully created, `inv1.log` contains the following error messages: Oracle not available, Out of space, Compilation error.

- 20** At the sqlplus prompt, enter `connect / as sysdba` to login as DBA.
- 21** At the sqlplus prompt, enter `shutdown immediate;`
- 22** At the Oracle Server Manager prompt, enter `@user_specified_directory_path/zenworks/_start.sql` to start the Inventory database.

Manually Creating the Oracle Inventory Database Object

- 1** In ConsoleOne, right-click a location in the Novell eDirectory tree for the database object, then click *New*, click *Object*, click *ZENworks Database*, then click *OK*.
- 2** Type a name for the database object, then click *OK*.
- 3** Configure the database server options of the database object.

3a In ConsoleOne, right-click the database object (Inventory database *_server_name*), then click *Properties*, then click the *ZENworks Database* tab.

3b Select the database server object by using either of the following methods:

- ♦ If eDirectory is installed on the database server, then in the *Server DN* field, browse for and select the Server object of the server where the database is physically installed and running.

The server's IP address is automatically populated to the *Server IP Address or DNS Name* drop-down list. If the selected server object has more than one IP address, select the appropriate IP address.

IMPORTANT: Make sure that the DNS name of the database server configured for the database object is valid. If the DNS name is invalid, you must select an appropriate database server IP address in the Database object property page.

To clear the value set in the *Server DN* field, type the IP address of another database server or browse and select another server object.

- ♦ If eDirectory is not installed on the database server, then specify the server IP address or the DNS name in the *Server IP Address or DNS Name* field.

3c Specify the following values:

- ♦ **Database (Read-Write) User Name:** *MW_DBA*
- ♦ **Database (Read-Write) Password:** *novell*
- ♦ **Database (Read Only) User Name:** *MWO_READER*
- ♦ **Database (Read Only) Password:** *novell*
- ♦ **Database (Write Only) User Name:** *MWO_UPDATER*
- ♦ **Database (Write Only) Password:** *novell*

IMPORTANT: All Inventory components use the username and password the configured in the database object. By default, “novell” is the password for all options. But you can change it in the database, and update the same here.

3d Click *Apply*.

3e To configure the JDBC Driver properties, click the *JDBC Driver Information* tab.

3f Select *Oracle*, then click the *Default Settings* button.

This populates the fields with default JDBC driver information.

The database settings for Oracle are:

- ♦ **Driver:** *oracle.jdbc.driver.OracleDriver*
- ♦ **Protocol:** *jdbc:*
- ♦ **SubProtocol:** *oracle:*
- ♦ **SubName:** *thin:@*
- ♦ **Port:** *1521*
- ♦ **Flags:** This field is not applicable for Oracle
- ♦ **Database Service Name:** *orcl* (The value for the SID is the same as assigned for the database instance.)

3g Click *Apply*, then click *Close*.

Configuring and Running Multiple Oracle Database Instances on a Windows Server

Make sure that the following prerequisites are met:

- The ZENworks supported Oracle version must be installed on the Windows Inventory server.
- To maintain the Inventory database in Oracle, Server Inventory requires that you have a minimum of 25 user licenses.
- You have already set up the Inventory database.

To configure and run Oracle instances:

- 1** At the database server, from the desktop *Start* menu, click *Programs*, click *Oracle*, click *Database Administration*, then click *Oracle Database Configuration Assistant*.
- 2** Click *Create a Database*, click *Next*, click *Typical*, click *Next*, click *Copy Existing Database Files from the CD*, then click *Next*.
- 3** Enter the following details:
 - ♦ **Global Database Alias** *mgmtdb.your_windows_nt/2000_name*
 - ♦ **SID:** By default, the value is *mgmtdb*.

4 Click *Finish*.

This process takes a significant amount of time and creates the Oracle database.

Make sure that the OracleServiceMGMTDB service is created and started.

5 Load the Inventory database.

From the desktop menu, click *Start*, click *Run*, then click *SQLPLUS* to run the Oracle Server Manager.

Enter the following commands:

```
set instance mgmtdb
connect internal/password_for_administrator
```

13.2.3 Setting Up the MS SQL Server 2000 or MS SQL Server 2005 Inventory Database

This section provides information on the following topics:

- ♦ “Configuring the MS SQL Server 2000 Inventory Database” on page 510
- ♦ “Configuring the MS SQL Server 2005 Inventory Database” on page 513
- ♦ “Manually Creating the Inventory Database Object for MS SQL 2000 or MS SQL 2005” on page 516
- ♦ “Connecting the Inventory Server and ConsoleOne to the MS SQL 2000 or MS SQL 2005 Inventory Database” on page 518

Configuring the MS SQL Server 2000 Inventory Database

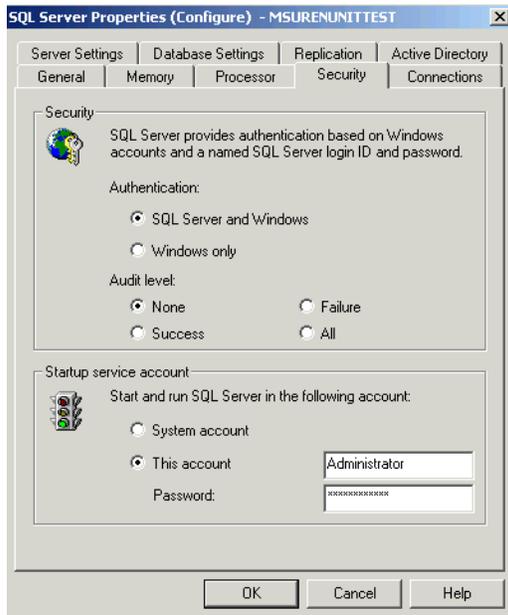
Prerequisites for configuring the database include the following:

- Microsoft SQL Server 2000 installed on the Windows server.
- Minimum free disk space of 50 MB to extract the `p1mssqlinvdb.zip` file.
- Make sure that you have sufficient disk space to store the inventory information on the server that has the Inventory database.

To configure the Inventory database for MS SQL Server 2000:

- 1 Copy the `p1mssqlinvdb.zip` file from the *ZENworks 7 Companion 2* CD\database\mssql directory to *path_of_inventory_database_directory_on_the_database_server*.
- 2 Extract `p1mssqlinvdb.zip`.
- 3 From the MS SQL server desktop *Start* menu, click *Programs*, click *Microsoft SQL Server*, then click *Enterprise Manager*.
- 4 In the SQL Server Enterprise Manager, browse to Console Root/Microsoft SQL Servers/SQL Server Group/*machine_name_running_Inventory_database*.
- 5 Right -click *machine_name_running_Inventory_database*, then click *Properties*.

- 6 In the SQL Server Properties dialog box, click the *Security* tab and make sure that the authentication is set to SQL Server and Windows.



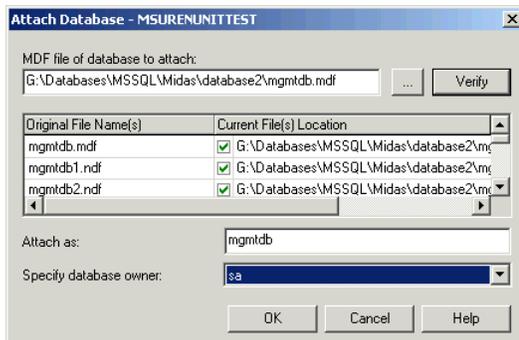
- 7 Click *OK*.

- 8 Browse to *machine_name_running_Inventory_database/Databases* and right-click *Databases*, click *All Tasks*, then double-click *Attach Database*.

- 9 In the Attach Database dialog box, do the following:

- 9a Click the *Browse* button to browse to and select *mgmt.db.mdf* as the *.mdf* database file to be attached.
- 9b Make sure that the value of the *Attach As* field is *mgmtdb*.
- 9c Select *sa* from the *Specify Database Owner* drop-down list.
- 9d Click *OK*.

The ZENworks Inventory database (*mgmtdb*) is attached to the *Databases* server group.



- 10 Select *mgmtdb*, then click the *Tools* menu, then click *SQL Query Analyzer*.

- 11 In the SQL Query Analyzer, do the following:

- 11a Make sure that *mgmtdb* is selected in the drop-down list.

11b Click *File*, then click *Open*.

11c Select the `createloginnames.sql` query file from *ZENworks 7 Companion 2* CD\database\mssql directory.

11d Click *Query*, then click *Execute*.

On successful execution, the following message is displayed in the Message pane:

```
New Login Created
```

11e Login as MW_DBA in the SQL Query Analyzer and execute the following drop trigger sqls:

```
drop trigger cim.x$cim$component
go
drop trigger cim.x$cim$dependency
go
drop trigger managewise.x$managewise$designates
go
drop trigger managewise.x$managewise$currentloginuser
go
drop trigger managewise.x$managewise$lastloginuser
go
drop trigger cim.x$cim$installedsoftwareelement
go
```

During the execution of the drop trigger sqls, the following error message might be displayed on the console, “Cannot drop the trigger '*trigger_name*', because it does not exist or you do not have permission”. Ignore the error message.

12 (Optional) Add non-English enumerated (enum) values for certain Inventory attributes into the Inventory database.

You need to add the non-English enumerated values so the Inventory ConsoleOne utilities such as Inventory Report can display the enum value for the inventory attributes in internationalized versions. The non-English enum values must be available in English version of the product so that the rolled-up inventory information from non-English sites can be properly captured at the high-level servers where only English versions are installed.

For more information about the list of attributes that contains enumerated values, see [Appendix M, “Enumeration Values,” on page 767](#).

To add the non-English enum values:

12a Specify the JDBC connection settings in the

```
zenworks_directory\inv\server\wminv\properties\
connection.prop file to connect to the MS SQL database.
```

You can do this by copying the template property settings for MS SQL specified in the comments section in the `connection.prop` file. Specify the IP address, port number, and Database SID in the JDBC URL string that matches your MS SQL server configuration.

12b At the server prompt, enter `AddEnums`

```
directory_name_containing_connection.prop
```

If your Inventory server is running on a Windows machine, run the above command from `zenworks_directory\inv\server\wminv\bin`.

After executing the command, the a message indicating that the non-English enums have been successfully inserted is displayed on the console prompt.

- 13 Continue with “**Manually Creating the Inventory Database Object for MS SQL 2000 or MS SQL 2005**” on page 516.

WARNING: Do not rename the mgmtdb database because it is set as the default database for the user account at login.

Configuring the MS SQL Server 2005 Inventory Database

NOTE: Information about MS SQL 2005 is applicable only for ZENworks 7 with Support Pack 1

Prerequisites for configuring the database include the following:

- Microsoft SQL Server 2005 installed on the Windows server.
- Minimum free disk space of 50 MB to extract the `p1mssqlinldb.zip` file.
- Make sure that you have sufficient disk space to store the inventory information on the server that has the Inventory database.

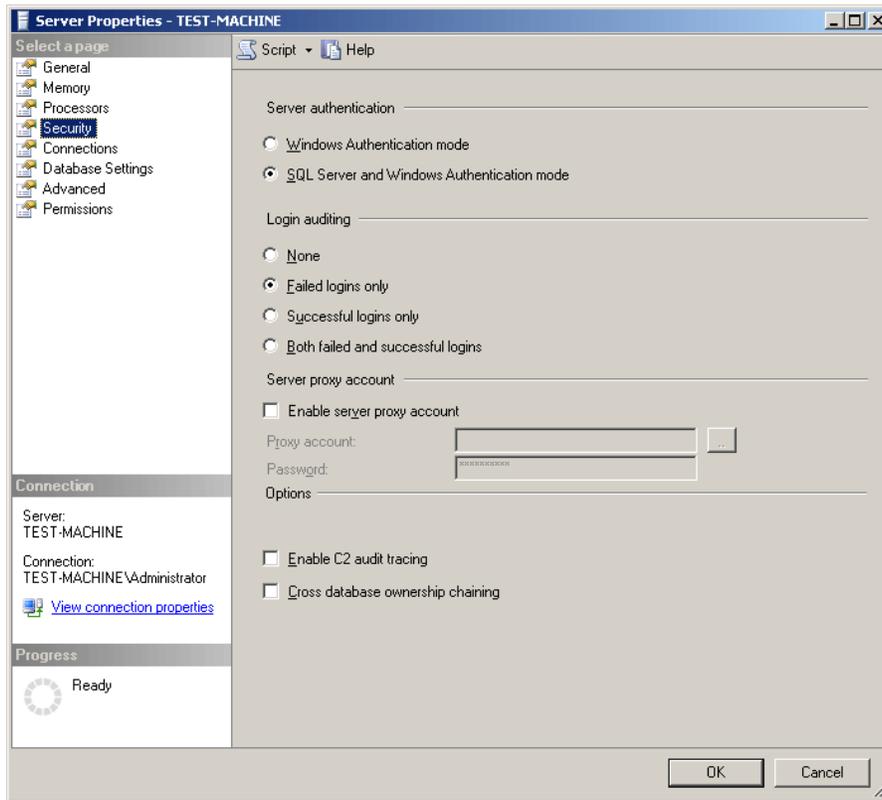
To configure the Inventory database for MS SQL Server 2005:

- 1 Copy the `p1mssqlinldb.zip` file from the *ZENworks 7 Companion 2* CD\`database\mssql` directory to *path_of_inventory_database_directory_on_the_database_server*.
- 2 Extract `p1mssqlinldb.zip`.
- 3 From the MS SQL server desktop *Start* menu, click *Programs*, click *Microsoft SQL Server 2005*, then click *SQL Server Management Studio*.
- 4 In the SQL Server Management Studio, connect to the Database Engine of the MSSQL Server by clicking on *File->Connect Object Explorer* and setting the properties.



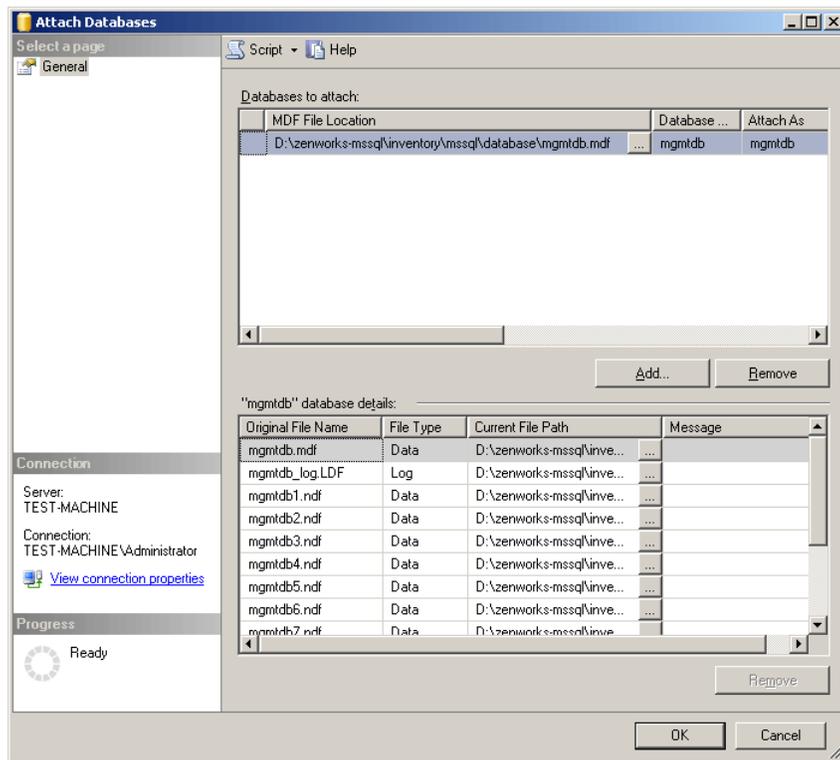
- 5 Browse to *machine_name_running_Inventory_database* in the Object Explorer, and right-click *machine_name_running_Inventory_database*, then click *Properties*.

- 6 In the SQL Server Properties dialog box, click the Security tab and make sure that the authentication is set to SQL Server and Windows.



- 7 Click OK.
- 8 Browse to *machine_name_running_Inventory_database/Database*, and right-click Databases, click All Tasks, then double-click Attach Database.
- 9 In the Attach Database dialog box, do the following:
 - 9a Click the *Add* button to browse to and select *mgmt.db.mdf* as the *.mdf* database file to be attached.
 - 9b Make sure that the value of the *Attach As* field is *mgmt.db*.
 - 9c Click *OK*.

The ZENworks Inventory database (mgmtdb) is attached to the Databases server group.



10 Right-click *mgmtdb*, then select *New Query*.

11 In the SQL Query Analyzer, do the following:

11a Select the `createloginnames.sql` query file from *ZENworks 7 Companion 2* CD\database\mssql directory. Either drag and drop it on the opened query window or copy the contents of the sql file to the query window.

11b Click *Execute*.

On successful execution, the following message is displayed in the Message pane:

New Login Created

11c Login as MW_DBA in the SQL Query Analyzer and execute the following drop trigger sqls:

```
drop trigger cim.x$cim$component
go
drop trigger cim.x$cim$dependency
go
drop trigger managewise.x$managewise$designates
go
drop trigger managewise.x$managewise$currentloginuser
go
drop trigger managewise.x$managewise$lastloginuser
go
drop trigger cim.x$cim$installedsoftwareelement
go
```

During the execution of the drop trigger sqls, the following error message might be displayed on the console, “Cannot drop the trigger '*trigger_name*', because it does not exist or you do not have permission”. Ignore the error message.

- 12 Continue with “[Manually Creating the Inventory Database Object for MS SQL 2000 or MS SQL 2005](#)” on page 516.

WARNING: Do not rename the mgmtdb database because it is set as the default database for the user account at login.

Manually Creating the Inventory Database Object for MS SQL 2000 or MS SQL 2005

NOTE: Information about MS SQL 2005 is applicable only for ZENworks 7 with Support Pack 1

- 1 In ConsoleOne, right-click a location in the Novell eDirectory tree for the database object, then click *New*, click *Object*, click *ZENworks Database*, then click *OK*.
- 2 Type a name for the database object, then click *OK*.
- 3 Configure the database server options of the database object.
 - 3a In ConsoleOne, right-click the database object (Inventory database_*server_name*), then click *Properties*, then click the *ZENworks Database* tab.

- 3b Select the database server object by using either of the following methods:

- ♦ If eDirectory is installed on the database server, then in the *Server DN* field, browse for and select the Server object of the server where the database is physically installed and running.

The server's IP address is automatically populated to the *Server IP Address or DNS Name* drop-down list. If the selected server object has more than one IP address, select the appropriate IP address.

IMPORTANT: Make sure that the DNS name of the database server configured for the database object is valid. If the DNS name is invalid, you must select an appropriate database server IP address in the Database object property page.

To clear the value set in the *Server DN* field, type the IP address of another database server or browse and select another server object.

- ♦ If eDirectory is not installed on the database server, then specify the server IP address or the DNS name in the *Server IP Address or DNS Name* field.

- 3c Specify the following values:

- ♦ **Database (Read-Write) User Name:** *MW_DBA*
- ♦ **Database (Read-Write) Password:** *novell*
- ♦ **Database (Read Only) User Name:** *MWM_READER*
- ♦ **Database (Read Only) Password:** *novell*
- ♦ **Database (Write Only) User Name:** *MWM_UPDATER*
- ♦ **Database (Write Only) Password:** *novell*

IMPORTANT: All Inventory components use the username and password the configured in the database object. By default, “novell” is the password for all options. But you can change it in the database, and update the same here.

3d Click *Apply*.

3e To configure the JDBC Driver properties, click the *JDBC Driver Information* tab.

3f If you have installed ZENworks 7 Desktop Management, select *MSSQL*, then click the *Default Settings* button.

This populates the fields with default JDBC driver information.

The database settings for MS SQL 2000 are:

- ♦ **Driver:** *com.microsoft.jdbc.sqlserver.SQLServerDriver*
- ♦ **Protocol:** *jdbc:*
- ♦ **SubProtocol:** *microsoft:*
- ♦ **SubName:** *sqlserver://*
- ♦ **Port:** *1433*
- ♦ **Flags:** This field is not applicable for MS SQL 2000.
- ♦ **Database Service Name:** This field is not applicable for MS SQL 2000.

3g If you have installed ZENworks 7 Desktop Management and MS SQL 2000, select *MSSQL (2000)*, then click the *Default Settings* button.

This populates the fields with default JDBC driver information.

The database settings for MS SQL 2000 are:

- ♦ **Driver:** *com.microsoft.jdbc.sqlserver.SQLServerDriver*
- ♦ **Protocol:** *jdbc:*
- ♦ **SubProtocol:** *microsoft:*
- ♦ **SubName:** *sqlserver://*
- ♦ **Port:** *1433*
- ♦ **Flags:** This field is not applicable for MS SQL 2000.
- ♦ **Database Service Name:** This field is not applicable for MS SQL 2000.

3h If you have installed ZENworks 7 Desktop Management with Support Pack 1 and MS SQL 2005, select *MSSQL (2005)*, then click the *Default Settings* button.

This populates the fields with default JDBC driver information.

The database settings for MS SQL 2005 are:

- ♦ **Driver:** *com.microsoft.jdbc.sqlserver.SQLServerDriver*
- ♦ **Protocol:** *jdbc:*
- ♦ **SubProtocol:** This field is not applicable for MS SQL 2005.
- ♦ **SubName:** *sqlserver://*
- ♦ **Port:** *1433*
- ♦ **Flags:** This field is not applicable for MS SQL 2005.
- ♦ **Database Service Name:** This field is not applicable for MS SQL 2005.

3i Click *Apply*, then click *Close*.

- 4 Continue with “[Connecting the Inventory Server and ConsoleOne to the MS SQL 2000 or MS SQL 2005 Inventory Database](#)” on page 518.

Connecting the Inventory Server and ConsoleOne to the MS SQL 2000 or MS SQL 2005 Inventory Database

NOTE: Information about MS SQL 2005 is applicable only for ZENworks 7 with Support Pack 1

The Inventory server components and the ConsoleOne use Microsoft JDBC driver to connect to the Inventory database on MS SQL 2000 or MS SQL 2005. You must install and configure Microsoft SQL Server 2000 or Microsoft SQL Server 2005 driver for JDBC driver with the Inventory system on MS SQL 2000 or MS SQL 2005 respectively.

To configure the Microsoft SQL Server 2000 or the Microsoft SQL Server 2005 driver for JDBC to access the Inventory database running on MS SQL 2000 or MS SQL 2005 respectively:

- 1 To configure the Microsoft SQL Server 2000 driver, download the Windows English version of Microsoft JDBC driver from the [Microsoft SQL Server Web site \(http://www.microsoft.com/downloads/details.aspx?FamilyID=9f1874b6-f8e1-4bd6-947c-0fc5bf05bf71&DisplayLang=en\)](http://www.microsoft.com/downloads/details.aspx?FamilyID=9f1874b6-f8e1-4bd6-947c-0fc5bf05bf71&DisplayLang=en).

NOTE: Skip to Step 5, if you have MS SQL 2005 as the database.

- 2 For Microsoft SQL Server 2000 driver, do the following on a Windows Inventory server and then skip to **Step 3**:
 - 2a Install the driver.
 - 2b Copy the `msbase.jar`, `msutil.jar`, and `mssqlserver.jar` files to `inventory_server_installation_directory\inv\server\lib` directory.
- 3 For Microsoft SQL Server 2000 driver, on the machine running ConsoleOne with ZENworks 7 Inventory snap-ins, copy the `msbase.jar`, `msutil.jar` and `mssqlserver.jar` files to the `consoleone_installation_directory\lib\zen` directory.
- 4 In ConsoleOne, create a database object in the same container where Inventory server is installed.
 - 4a Right-click the container.
 - 4b Click *New*, click *Object*, select *ZENworks Database* from the list of objects, then click *OK*.
 - 4c Enter a name for the database object, then click *OK*.
- 5 Configure the Database server options of the Database object.
 - 5a In ConsoleOne, right-click the database object, click *Properties*, then click the *ZENworks Database* tab.
 - 5b Select the database server object using any of the following methods:
 - ♦ If eDirectory is installed on the database server, in the *Server DN* field, browse for and select the Server object for the server where the database is physically installed and running.

The server’s IP address is automatically populated to the *Server IP Address or DNS Name* drop-down list. If the selected server object has more than one IP address, select the appropriate IP address.

IMPORTANT: Make sure that the DNS name of the database server configured for the database object is valid. If the DNS name is invalid, you must select an appropriate database server IP address in the Database object property page.

To clear the value set in the *Server DN* field, type the IP address of another database server or browse and select another server object.

- ♦ If eDirectory is not installed on the database server, then enter the server's IP address or the DNS name in the *Server IP Address or DNS Name* field.

5c Type the values for the following options:

- ♦ **Database (Read-Write) User Name:** *MW_DBA*
- ♦ **Database (Read-Write) Password:** *novell*
- ♦ **Database (Read Only) User Name:** *MWM_READER*
- ♦ **Database (Read Only) Password:** *novell*
- ♦ **Database (Write Only) User Name:** *MWM_UPDATER*
- ♦ **Database (Write Only) Password:** *novell*

5d Click *Apply*.

5e To configure the JDBC Driver properties, click the *JDBC Driver Information* tab.

5f For MS SQL 2000 , select *MS SQL* or *MS SQL (2000)* in case of ZENworks 7, then click *Default Settings*.

This populates the fields with default JDBC driver information.

Modify the database settings based on the configuration of your MS SQL Server. The database settings for MS SQL are:

- ♦ **Driver:** *com.microsoft.jdbc.sqlserver.SQLServerDriver*
- ♦ **Protocol:** *jdbc:*
- ♦ **SubProtocol:** *microsoft:*
- ♦ **SubName:** *sqlserver://*
- ♦ **Port:** *1433*
- ♦ **Flags:** Not applicable for MS SQL 2000
- ♦ **Database Service Name:** Not applicable for MS SQL 2000

5g For MS SQL 2005, select *MS SQL (2005)*, then click *Default Settings*.

This populates the fields with default JDBC driver information.

Modify the database settings based on the configuration of your MS SQL Server. The database settings for MS SQL are:

- ♦ **Driver:** *com.microsoft.jdbc.sqlserver.SQLServerDriver*
- ♦ **Protocol:** *jdbc:*
- ♦ **SubProtocol:** Not applicable for MS SQL 20005
- ♦ **SubName:** *sqlserver://*
- ♦ **Port:** *1433*
- ♦ **Flags:** Not applicable for MS SQL 2005
- ♦ **Database Service Name:** Not applicable for MS SQL 2005

5h Click *Apply*, then click *Close*.

For more information on Performance tips, see [Section J.1, “Database Parameter Tuning Tips,”](#) on page 717.

13.3 Configuring the Inventory Service Object

The Inventory Service object settings configure the scanning for the associated inventoried servers.

To configure the Inventory Service Object:

- 1 In ConsoleOne, right-click the Inventory Service object (Inventory Service_ *server_name*), then click *Properties* to display the Inventory Service Object Properties page.

- 2 Modify the following settings:

Inventory Server Role: Based on the Inventory servers that you have deployed for scanning inventory, you must specify the role of the Inventory server. See [Section 12.4, “Understanding the Inventory Server Roles,”](#) on page 448.

Based on the new role you select, you will see a list of actions to be followed. For example, if you change the role of the Root Server to Root Server with Inventoried Servers, you must configure the Server Inventory policy for the inventoried servers that you have attached.

Similarly, to change the role to any other server, follow the actions to make the new role change effective. For more information, see [Section 13.1.5, “Changing the Role of the Inventory Server,”](#) on page 483.

Plan the change of roles carefully because these changes impact the existing inventory deployment.

Discard Scan Data Time: Select the date and time. Any scan data files (.zip files) that have scan information collected before the Discard Scan Data Time that you specify in the Inventory Service Object property page are discarded.

Scan Directory Path: Select the name of the volume on the Inventory server where you want to store the scan data files.

The Scan directory (`scandir`) path is the location on the Inventory server that stores the scan data files. The format of the Scan directory path is as follows:
inventory_server_name\volume_of_the_server_directory.

For a NetWare server, you cannot modify the Inventory server name specified in the Scan directory path. To modify the directory name, click *Browse* and select an existing directory.

For a Windows server, you cannot modify the Inventory server name specified in the Scan directory path. To modify the directory name, you must manually type it.

Enable Scan of Machines: Select this option to specify hardware and software scanning of the inventoried servers associated with this Inventory Service object. The scanners collect inventory information only if this option is enabled. By default, the scanners collect only hardware information for the inventoried servers.

- 3 To configure the software dictionary rules, click the *Software Inventory Configuration* tab. For more information on how to configure the software dictionary rules, see [Section 16.3, “Customizing the Software Inventory Information To Be Scanned For the ZENworks 7 Inventoried Servers,”](#) on page 586.

- 4 Click *OK*.

NOTE: If you are modifying the Inventory policies or configuring the objects, always stop the Inventory services. Configure the policies and properties of the objects. Restart the Inventory services again. For more information on how to start the inventory service, see [Section 13.1.4, “Starting and Stopping the Inventory Service,”](#) on page 482.

13.4 Configuring the Database Location Policy

The Database Location policy contains the location of the Inventory database. You can associate the Database object with a container under which the Inventory Service object is located through using the Service Location Package or with an Inventory server through using the Server Package.

NOTE: If you configure the Service Location Package and the Server Package, the Server Package settings override the Service Location Package settings.

To associate the Database object with a container under which the Inventory Service object is located:

- 1 In ConsoleOne, right-click the Service Location Package, then click *Properties* to display the Policies page.
- 2 Select the check box under the *Enabled* column for the ZENworks Database policy.
- 3 Click *Properties* to display the *Inventory Management* page.
- 4 Browse to the DN of the Inventory Database object (Inventory database_ *server_name*), then click *OK*.

For a Sybase database, the database object is automatically created during the Server Inventory installation unless you are installing on a Windows server without eDirectory installed. To manually create the database object, see [“Manually Creating the Sybase Inventory Database Object”](#) on page 494.

For an Oracle database, you must create the database object and configure the object. For more information, see [“Setting Up the Oracle Inventory Database”](#) on page 500.

For an MS SQL database, you must configure the database object. For more information, see [“Setting Up the MS SQL Server 2000 or MS SQL Server 2005 Inventory Database”](#) on page 510.

- 5 Click *OK*.
- 6 Click the *Associations* tab, then click *Add*.
- 7 Browse to select the container under which the Inventory Service object is located, then click *OK*.
- 8 Click *Apply*, then click *Close*.

To associate the Database object with an Inventory server:

- 1 In ConsoleOne, right-click the Server Package, click *Properties* to display the *Policies* page.
- 2 Select the check box under the *Enabled* column for the ZENworks Database policy.
- 3 Click *Properties* to display the *Inventory Management* page.
- 4 Browse to the DN of the Inventory Database object (Inventory database_ *server_name*), then click *OK*.

For a Sybase database, the database object is automatically created during the Server Inventory installation unless you are installing on a Windows server without eDirectory installed. To manually create the database object, see [“Manually Creating the Sybase Inventory Database Object” on page 494](#).

For an Oracle database, you must create the database object and configure the object. For more information, see [“Setting Up the Oracle Inventory Database” on page 500](#).

For an MS SQL database, you must create the database object and configure the object. For more information, see [“Setting Up the MS SQL Server 2000 or MS SQL Server 2005 Inventory Database” on page 510](#).

- 5 Click *OK*.
- 6 Click the *Associations* tab, then click *Add*.
- 7 Browse to select an Inventory server object, then click *OK*.
- 8 Click *Apply*, then click *Close*.

NOTE: If you are modifying the Inventory policies or configuring the objects, always stop the Inventory services. Configure the policies and properties of the objects. Restart the Inventory services again. For more information on how to start the inventory service, see [Section 13.1.4, “Starting and Stopping the Inventory Service,” on page 482](#).

13.5 Configuring the Server Inventory Policy

The Server Inventory policy contains the IP address or the DNS name of the Inventory server to which the inventory information is sent. This policy also contains the inventory scanning schedule for the associated inventoried server. You must configure the Server Inventory policy for each inventoried server.

To configure the Server Inventory policy package:

- 1 In ConsoleOne, right-click the Distributed Server Package, then click *Properties* to display the *Policies* page.
- 2 Click the *Policies* tab, then click *NetWare* or *Windows* from the drop-down list, depending on the operating system of the inventoried server.
- 3 Select the check box under the *Enabled* column for the Server Inventory policy.
- 4 Click *Properties* to display the Server Inventory Policy page.
- 5 In the General tab, configure the following settings:
 - 5a Browse to select the DN of the Inventory Service object (Inventory Service_ *server_name*).

This setting specifies that the scanner will send the server inventory information to this Inventory server.
 - 5b Select the DNS name or the IP address of the Inventory server.
 - 5c If the scan is to an Inventory server that is across the firewall, specify the IP address and the port number of the proxy server.
- 6 (Optional) Customize the Inventory scan:
 - 6a To customize the hardware scan for the Windows inventoried servers, click the *Hardware Scan* tab and configure the following settings:

Enable DMI Scan: Includes DMI scanning of Windows inventoried servers.

Enable WMI Scan: Includes WMI scanning of Windows inventoried servers.

- 6b** To customize the software scan for the NetWare or Windows inventoried servers on which Novell ZENworks for Servers 3.0 or ZENworks for Servers 3.0.2 is installed, click the Software Scan tab and configure the following settings. For more information, see [Section 16.4, “Customizing the Software Inventory Information To Be Scanned For ZENworks for Servers 3.x Inventoried Servers,”](#) on page 632.

IMPORTANT: Do not configure these settings for the inventoried servers that have ZENworks 7 Server Management installed. To customize software scanning for servers having ZENworks 7 Server Management installed, see [Section 16.3, “Customizing the Software Inventory Information To Be Scanned For the ZENworks 7 Inventoried Servers,”](#) on page 586.

Enable Software Scan: Enables software scanning for the inventoried servers associated with the Inventory policy. The scan program collects software information for the inventoried servers and stores it in the Inventory database.

Custom Scan Editor: Enables you to customize the list of application details to scan for at the Windows inventoried servers. The Inventory scanner scans for the details of the applications listed in the Custom Scan Editor.

For example, specify the following details in the Custom Scan Editor: Vendor Name=Microsoft; Product Name=Microsoft Office; Product Version=10.0; FileName=winword.exe; File Size=1 MB. The Inventory scanner scans for the winword.exe file having a size of 1 MB on the inventoried servers. If the file is found, the scanner stores “Microsoft;Microsoft Office;10.0” for “winword.exe;1 MB” in the Inventory database.

Product Identification Number: Enables you to scan for the product identification number of the Microsoft applications installed on the Windows inventoried servers only

- 6c** Click the *Configuration Editor* tab; if required, modify the settings of the following .ini files.
- ♦ **Asset Information:** Use this file to scan for vendor-specific information from DMI. For more information on how to configure the Asset Information, see [“Scanning for Vendor-Specific Asset Information from DMI”](#) on page 582.
 - ♦ **Zippered Names:** Use this file to customize the hardware scanning of Jaz* and Zip* drives. For more information, see [“Customizing the Hardware Scanning Information of Jaz and Zip Drive Vendors”](#) on page 583.
 - ♦ **SWRules:** Configure the SWRules file for the Windows inventoried servers on which Novell ZENworks for Servers 3.0 or ZENworks for Servers 3.0.2 is installed. Do not configure the settings for inventoried servers that have ZENworks 7 Server Management installed.

The SWRules option customizes the software scanning information of vendors and products. For more information, see [Section 16.4, “Customizing the Software Inventory Information To Be Scanned For ZENworks for Servers 3.x Inventoried Servers,”](#) on page 632.
 - ♦ **HWRules:** Use this file to customize the information on nominal size of monitors. For more information on how to configure the HWRules .ini file, see [Section 16.2.3, “Customizing the Hardware Information for Monitor Size,”](#) on page 584.

- 7 Click the *Policy Schedule* tab.
- 8 Modify the schedule, click *Apply*, then click *Close*.
- 9 In the Distributed Server Package property page, click the *Distribution* tab, then click *Add*.
- 10 Browse to add the Distribution object, then click *OK*.
- 11 Click *Apply*, then click *Close*.
- 12 In ConsoleOne, right-click the Inventory Service object, click *Properties*, then click the *Inventory Service Object Properties* tab.
- 13 Ensure that the *Enable Scan of Machines* check box is selected, then click *OK*.
This setting ensures that scanning is selected for the inventoried servers associated with the selected Inventory server.

13.6 Configuring the Roll-Up Policy

The Roll-Up policy settings configure the selected Inventory server for roll-up of scan information. The settings in the Roll-Up policy identify the next-level Inventory server (DN of the Inventory Service object) for moving the inventory information from the selected Inventory server. These settings stored in eDirectory are associated with the Inventory Server object.

To configure the Roll-Up policy:

- 1 In ConsoleOne, right-click the Server Package, click *Properties*, click *Policies*, then select the appropriate suboption. If you want to this policy to be applied on all servers, select the General suboption.
- 2 Check the check box under the *Enabled* column for the Inventory Rollup Policy.
- 3 Click *Properties* to display the Roll-Up Policy page.
- 4 Browse to select the DN of the Inventory Service object (*Inventory Service_server_name*).

Destination Server Object: You must specify the DN of the Inventory Service object at the next level Inventory server for moving the inventory information from the selected Inventory server. The server that you specify must be another Intermediate Server, Intermediate Server with Database, Intermediate Server with Database and Inventoried Servers, Intermediate Server with Inventoried Servers, Root Server, or Root Server with Inventoried Servers.

NOTE: Ensure that the specified Inventory server is a different server because you cannot roll-up of information to the same Inventory server. Also, you cannot specify the lower-level Inventory server as the next-destination server for roll-up of information.

- 5 By default, the DNS name or the IP address (if a DNS name is not configured) of the next-level server is populated in the field. If the next-level server has multiple IP addresses, select the preferred address.

IMPORTANT: Ensure that the DNS name of the next-level server is valid. If the DNS name is invalid, you must select an appropriate server IP address.

- 6 If the roll-up is to an Inventory server that is across the firewall, specify the IP address or the DNS name and the port number of the proxy server.
- 7 Click *Apply*.
- 8 Click the *Roll-Up Policy* tab, then click *Roll-Up Schedule*.
- 9 Modify the settings for scheduling the roll-up time and click *Apply*.

When you schedule the roll-up of information in the Inventory policies, we recommend the roll-up frequency should be at least one day. It is likely that if the roll-up of inventory information is scheduled too frequently, for example less than one hour, there might be some performance degradation of the Inventory server.

- 10** (Conditional) If you have not yet associated the Server Package, you are prompted to associate it with an Inventory server or a container. The policy you configured and enabled earlier will not be in effect until you associate this policy package with an Inventory server or a container.

To associate the policy package:

10a Click the *Associations* tab, then click *Add*.

10b Browse for and select the Inventory server or the container that you want to associate the Roll-Up policy to.

10c Click *OK*, then click *OK*.

- 11** Click *Apply*, then click *Close*.

NOTE: If you are modifying the Inventory policies or configuring the objects except for the Roll-Up schedule, always stop the Inventory services. Configure the policies and properties of the objects. Restart the Inventory services again. For more information, see [Section 13.1.4, “Starting and Stopping the Inventory Service,” on page 482](#).

13.7 Configuring the Dictionary Update Policy

The Dictionary Update policy configures the Inventory server to receive the software dictionary updates from other Inventory servers. You must manually download the dictionary updates to at least one Inventory server in your network. This Inventory server can then act as the source of dictionary updates to other Inventory servers.

- 1** In ConsoleOne, right-click the Server Package, click *Properties*, click *Policies*, then select the appropriate suboption. If you want to this policy to be applied on all servers, select the *General* suboption.
- 2** Select the check box under the *Enabled* column for the Dictionary Update policy.
- 3** Click *Properties* to display the Dictionary Update Policy page.
- 4** Configure the following settings:
 - 4a** (Recommended) Select the *Use the Roll-Up Server as the Update Source* check box if you want the Dictionary Update Service to use the Inventory server configured in the Roll-Up policy as the source for dictionary updates.

If you select this check box, continue with [Step 9 on page 526](#). If you do not select this option, the Dictionary Update Service will use the following settings configured in this policy (Dictionary Update Policy); continue with [Step 4b on page 525](#).

NOTE: Do not select this option for a Standalone Server and a Root Server. You must manually configure the settings of the policy.

- 4b** In the *Source Service Object* field, browse to select the DN of the Inventory server, which provides the dictionary updates.
- 4c** Select the IP address or the DNS name of the Inventory server, which provides the dictionary updates.

- 4d** If the source Inventory server is across the firewall, specify the IP address or the DNS name and the port number of the XML proxy server.
- 4e** Click *Apply*.
- 5** Click the *Dictionary Update Policy* tab, then click *Dictionary Update Schedule*.
- 6** Configure the Dictionary Update Schedule page to establish the schedule for running the Dictionary Consumer.
- We recommend you to configure the Weekly schedule.
- 7** Click *Apply*.
- 8** (Conditional) If you have not yet associated the Server Package, you are prompted to associate it with an Inventory server or a container. The policy you configured and enabled earlier will not be in effect until you associate this policy package with an Inventory server or a container. To associate the policy package:
- 8a** Click the *Associations* tab, then click *Add*.
- 8b** Browse for and select the Inventory server or the container that you want to associate the Dictionary Update policy to.
- 8c** Click *OK*, then click *OK*.
- 9** Click *Apply*, then click *Close*.

NOTE: If you want to modify the Dictionary Update policy settings, you need not stop the Inventory services.

13.8 Setting Up Distribution of Dictionary

A software dictionary can be updated in the following ways:

- ♦ Manually download the dictionary from TID 10093255 in the [Novell Support Knowledgebase \(http://www.novell.com/support\)](http://www.novell.com/support) to each Inventory server.

NOTE: The dictionary is updated and published once every three months in this TID.

- ♦ Manually download the dictionary from the [Novell Support Knowledgebase \(http://www.novell.com/support\)](http://www.novell.com/support) to one Inventory server (preferably, the Root Server) and automatically distribute the dictionary to all servers in your setup by configuring the Dictionary Update policy.

An Inventory server can receive dictionary updates from any other Inventory server, irrespective of the server's role. The role of the Inventory server indicates whether the server receives the inventory information, stores the information into a local Inventory database, or rolls up the inventory information.

To update and distribute the software dictionary between Inventory servers:

- 1** Manually download the dictionary from TID 10093255 in the [Novell Support Knowledgebase \(http://www.novell.com/support\)](http://www.novell.com/support) and save it in the `zenworks_installation_directory\zenworks\inv\server\dictdir` directory on the Inventory server.
- 2** Configure the Dictionary Update policy. For more information see, [Section 13.7, “Configuring the Dictionary Update Policy,” on page 525.](#)

All Inventory servers have Dictionary Provider and Dictionary Consumer services that are automatically installed during the Server Inventory installation.

When an Inventory server is started, the Dictionary Consumer reads the Dictionary Update policy and contacts the Dictionary Provider (running on another Inventory server) specified in the policy.

Subsequently, the Dictionary Consumer checks for the dictionary updates based on the schedule set in the Dictionary Update policy. It compares the date of the dictionary file on the Dictionary Provider with the file that has been locally stored. If the file on the Dictionary Provider is newer, then the Dictionary Consumer downloads the file from the Dictionary Provider using XML-RPC as per the schedule.

The user-defined rules in the downloaded dictionary file are merged with the rules present in the local dictionary. If the merge yields a different set of rules from those locally present, the consolidated set of rules is written to the local dictionary. During the merge process, conflicts might arise, which are resolved on the basis of the following considerations:

- ◆ The rules in the downloaded dictionary always override the rules in the local dictionary.
- ◆ If a conflict arises between the software identifiers, the conflicting identifiers in the local dictionary are removed from the final (merged) dictionary.
- ◆ For a software dictionary rule, the final result is obtained by first writing the downloaded rules and then the local rules into the final dictionary; eliminating the duplicates during the process. This ensures that the downloaded software rules precede the local rules.

The following scenario illustrates the distribution of the software dictionary among the Inventory servers.

In this scenario, there is an Inventory tree consisting of one Root Server (R1), one Leaf Server (L1), and two Standalone servers (S1 and S2). L1 rolls up the inventory information to R1.

Follow the below procedure to update the software dictionary on all the Inventory servers.

1. Manually download the dictionary on R1 from TID 10093255 in the [Novell Support Knowledgebase](http://www.novell.com/support) (<http://www.novell.com/support>).
2. Create and configure a Dictionary Update policy by specifying R1 as the Dictionary Provider, and associate the policy to L1. For more information on how to configure the Dictionary Update policy, see [Section 13.7, “Configuring the Dictionary Update Policy,” on page 525](#).
3. For S1 and S2, you can either manually download the dictionary from TID 10093255 in the [Novell Support Knowledgebase](http://www.novell.com/support) (<http://www.novell.com/support>) or configure the Inventory servers to automatically receive the latest version of the dictionary from R1.

For S1 and S2 to automatically receive the latest version of the dictionary from R1, create and configure a Dictionary Update policy by specifying R1 as the Dictionary Provider, and associate the policy to S1 and S2. For more information on how to configure the Dictionary Update policy, see [Section 13.7, “Configuring the Dictionary Update Policy,” on page 525](#).

Understanding the Server Inventory Components

14

The following sections describe the Novell® ZENworks® 7 Server Inventory components and processes:

- ◆ [Section 14.1, “Understanding the Inventory Service Manager,” on page 529](#)
- ◆ [Section 14.2, “Understanding the Server Configuration Service,” on page 532](#)
- ◆ [Section 14.3, “Understanding the Inventory Scanner,” on page 532](#)
- ◆ [Section 14.4, “Understanding the Sender and Receiver,” on page 539](#)
- ◆ [Section 14.5, “Understanding the Selector,” on page 544](#)
- ◆ [Section 14.6, “Understanding the Storer,” on page 545](#)
- ◆ [Section 14.7, “Understanding the Dictionary Provider and Dictionary Consumer,” on page 546](#)
- ◆ [Section 14.8, “Understanding the Upgrade Service,” on page 546](#)
- ◆ [Section 14.9, “An Overview of the Inventory Components on the Inventory Server,” on page 546](#)
- ◆ [Section 14.10, “Understanding the Inventory Database,” on page 547](#)

14.1 Understanding the Inventory Service Manager

The Inventory Service Manager loads the inventory components on the Inventory server, based on the configuration parameters specified in the Inventory server properties file.

This section contains the following:

- ◆ [Section 14.1.1, “List of Services,” on page 529](#)
- ◆ [Section 14.1.2, “Services on NetWare Inventory Servers,” on page 530](#)
- ◆ [Section 14.1.3, “Services on Windows Inventory Servers,” on page 530](#)

14.1.1 List of Services

The Service Manager loads the following important services. You can obtain the list of services that the Service Manager loads from the property file in

```
inventory_server_installation_directory_or_volume\zenworks\inv\server\wminv\properties.
```

Server Configuration Service
Inventory Scheduler Service
Inventory Scheduler Service
Selector Service
Receiver Service
Sender Service

Storer Service
Scan Collector Service
Dictionary Consumer Service
Dictionary Provider Service

You can use these service names to list, start, and stop the respective services.

The Inventory Service Manager reads the server property file (`config.properties`) and the role-based property file in the `inventory_server_installation_directory_or_volume\zenworks\inv\server\wminv\properties` directory, and loads the required services and server components.

IMPORTANT: Do not modify the property files because the updates might fail to load the services or the Service Manager.

14.1.2 Services on NetWare Inventory Servers

On a NetWare[®] Inventory server, the installation program modifies the `autoexec.ncf` file located in `sys:\system` directory to load `startinv.ncf`. The `startinv.ncf` file located in the `sys:\system` brings up the Inventory Service Manager at Inventory server startup time.

You can start, stop, or list the services, if the Inventory Service Manager is already loaded.

- ◆ To check if the Inventory Service Manager is loaded, at the server prompt, enter `java -show`.

This displays the following message:

```
com.novell.zenworks.inventory.servercommon.ZENWorksInventoryServiceManager
```

- ◆ To start an Inventory service, enter `StartSer service_name` at the Inventory server prompt.

`service_name` refers to any of the listed services. Follow the service naming syntax when you modify the `service_name`.

For example, to start the Storer, enter `StartSer Storer`

- ◆ To stop an Inventory service, enter `StopSer service_name` at the Inventory server prompt.

`service_name` refers to any of the listed services. Follow the service naming syntax when you modify the `service_name`.

For example, to stop the Storer, enter `StopSer Storer`

- ◆ To list an Inventory service, enter `ListSer service_name` at the Inventory server prompt.
`service_name` refers to any of the listed services. Follow the service naming syntax when you modify the `service_name`.
- ◆ To list all Inventory services, enter `ListSer *` at the Inventory server prompt.

14.1.3 Services on Windows Inventory Servers

On Windows Inventory servers, the installation program creates the Service Manager as a service. During server startup, this Inventory Service Manager is loaded as a service.

You can start, stop, or list the services, if the Inventory Service Manager (ZENworks Inventory Service) is already loaded.

To start the Inventory service on the Windows 2000/2003 Inventory server:

- 1 In the Control Panel, double-click *Administrative Tools*.
- 2 Double-click *Services*.
- 3 Select *Novell Inventory Service*, then click *Start*.

To stop the Inventory service on the Windows 2000/2003 Inventory server:

- 1 In the Control Panel, double-click *Administrative Tools*.
- 2 Double-click *Services*.
- 3 Select *Novell Inventory Service*, then click *Stop*.

To start an Inventory service:

- 1 Go to the `installation_directory\inv\server\wminv\bin` directory.
- 2 At the prompt, enter `StartSer service_name`.
where `service_name` refers to an Inventory service. Follow the service naming syntax when you modify the `service_name`.
For example, to start the Storer, enter `StartSer Storer`

To stop an Inventory service:

- 1 Go to the `installation_directory\inv\server\wminv\bin` directory.
- 2 At the prompt, enter `StopSer service_name`.
where `service_name` refers to an Inventory service. Follow the service naming syntax when you modify the `service_name`.
For example, to stop the Storer, enter `StopSer Storer`

To stop all Inventory services (ZENworks Inventory Service), use the Windows services from the desktop menu.

To list an Inventory service:

- 1 Go to the `installation_directory\inv\server\wminv\bin`.
- 2 At the prompt, enter `ListSer [-verbose] service_name`.
where `service_name` refers to an Inventory service.

Follow the service naming syntax when you modify the `service_name`.

To refer to all services, use the asterisk (*) wildcard character within double quotes `"*`". This wildcard character can be used with `ListSer` parameters. For example, to list all Inventory services, enter `ListSer "*"`.

14.2 Understanding the Server Configuration Service

The Server Configuration Service performs the following tasks:

- ◆ Reads the policy information from the Novell eDirectory™ and passes it to other Inventory components.
- ◆ Validates the policies to ensure that the policies are correctly configured.
- ◆ Validates the Inventory database version.

14.3 Understanding the Inventory Scanner

ZENworks 7 Server Management uses the Inventory scanner to collect hardware and software information from NetWare and Windows inventoried servers.

The scanner collects hardware details such as: floppy disk drive, hard disk drive, BIOS, bus, mouse, keyboard, display adapters, network adapter cards, modems, Jaz drives, Zip drives, sound cards, memory cards, serial ports, parallel ports, processors, and modems. The software scanning includes checking for applications on the inventoried servers and reporting the information about the scanned software, such as the vendor name, the product name and version, etc.

The following sections contain detailed information about the Inventory scanners:

- ◆ [Section 14.3.1, “Inventory Scanning Process,” on page 532](#)
- ◆ [Section 14.3.2, “Scanning for the NetWare Inventoried Servers,” on page 533](#)
- ◆ [Section 14.3.3, “Scanning for the Windows Inventoried Servers,” on page 535](#)

You can customize the hardware information and the software information to be scanned. For more information, see [Section 16.2, “Customizing the Hardware Inventory Information To Be Scanned,” on page 582](#) and [Section 16.3, “Customizing the Software Inventory Information To Be Scanned For the ZENworks 7 Inventoried Servers,” on page 586](#).

14.3.1 Inventory Scanning Process

1. Subscriber must be installed and configured on the inventoried servers.
2. The Server Inventory policy lets you configure the scanning schedule based on which the policy engine schedules and enforces scanning at the inventoried servers.
3. The Inventory scanner checks whether an updated dictionary is available at its Inventory server and downloads the updated dictionary.
4. The Policy Enforcer first reads the inventory settings configured in the Inventory Service object and the Server Inventory policy, and then launches the Inventory scanner.
5. The scanner scans for the hardware and software information.
6. The scan information collected by the scanners is stored as scan data files (.str). The files are sent to the Inventory server.

14.3.2 Scanning for the NetWare Inventoried Servers

The Inventory scanner scans for the hardware and software inventory information on the NetWare inventoried servers. For more information, review the following sections:

- ♦ [“Scanning for the Hardware Inventory Information” on page 533](#)
- ♦ [“Scanning for the Software Inventory Information” on page 533](#)

Scanning for the Hardware Inventory Information

Following are the sources on the NetWare inventoried servers from where the hardware inventory information is scanned:

- ♦ [“Simple Network Management Protocol \(SNMP\)” on page 533](#)
- ♦ [“SMBIOS” on page 533](#)
- ♦ [“Probe” on page 533](#)

For more information about the hardware information collected by the Inventory scanner, see [Appendix K, “Hardware Information Collected by the Inventory Scanners,” on page 729](#).

Simple Network Management Protocol (SNMP)

The scanners collect certain hardware (device) and network information based on SNMP. Additionally, the scanner also uses SNMP to report software installed and registered in `products.dat`. The scanner uses SNMP v2.0 and the services of `hostmib.nlm`, `ipxrtr.nlm`, and `ipxrtrnm.nlm`.

SMBIOS

The Inventory scanners use the SMBIOS information embedded into the BIOS of most hardware to report BIOS version, BIOS release date, Manufacturer, asset tag, product name, serial number, processor information, cache information, system slots information, port information, video adapter name, sound card name, and so on. The Inventory scanner reads information from SMBIOS with the help of `invaidd.nlm`.

Probe

Probe is a special built-in algorithm in the Inventory scanner, which is used to collect hardware information.

Scanning for the Software Inventory Information

The Inventory scanner scans the following software inventory information on the NetWare inventoried servers:

- ♦ [“Installed Software Information” on page 534](#)
- ♦ [“Disk Usage” on page 534](#)
- ♦ [“File Information” on page 535](#)
- ♦ [“AntiVirus Definition Files” on page 535](#)

Installed Software Information

The scanner collects software information from the following sources on the inventoried servers: Microsoft* Installer (MSI), Add-Remove Programs, Dictionary-based scan, and Probe.

Products.Dat: Includes software that are installed on the NetWare inventoried servers.

Dictionary-based scan: Includes software that are collected based on the software dictionary rules. For more information, see [Section 16.3, “Customizing the Software Inventory Information To Be Scanned For the ZENworks 7 Inventoried Servers,”](#) on page 586.

Probe: Probe is a special built-in algorithm in the Inventory scanner, which is used to collect software information about ZENworks Suite and its installed components.

[Table 14-1](#) shows the software information collected by the scanner from the various sources:

Table 14-1 *Software Information collected by the NetWare Inventory Scanner*

| Scanned Attributes | Product.Dat | Dictionary-based scan | Probe |
|--------------------------|-------------|-----------------------|-------|
| Product Name | Yes | Yes | Yes |
| Vendor Name | No | Yes | Yes |
| Product Version | Yes | Yes | Yes |
| Product Identifier | No | No | No |
| Product Install Location | No | Yes | Yes |
| Category | No | Yes | No |
| Description | No | Yes | No |
| Help Link | No | No | No |
| MSI Package GUID | No | No | Yes |
| Display/Internal Version | No | Yes | Yes |
| Language | No | No | Yes |
| UnInstall String | No | No | No |
| Installation Source | No | No | No |
| Display Name | No | No | Yes |
| Support Pack | No | No | Yes |
| Product Edition | No | No | Yes |
| Last Execution Time | No | No | No |
| Usage Count | No | No | No |

Disk Usage

The scanner collects the total disk usage information for the file extensions that are configured in the Software Dictionary editor. For more information, see [Section 16.3, “Customizing the Software Inventory Information To Be Scanned For the ZENworks 7 Inventoried Servers,”](#) on page 586.

File Information

The scanner reports the following information for the files that match with the dictionary entries, and the files that belong to the unidentified software list, which is configured using the software dictionary rules. For more information, see [Section 16.3.10, “Configuring the Software Dictionary Rules,” on page 596](#)

The scanner reports the following file attributes: FileName, FileSize, LastModifiedTime, InternalName, FileVersion, ProductName, ProductVersion, CompanyName, Language, DirectoryPath, and SoftwareDictionaryID.

AntiVirus Definition Files

The scanner collects information about the latest virus definition date and version of McAfee Netshield* 4.6.x installed on the inventoried servers.

14.3.3 Scanning for the Windows Inventoried Servers

The Inventory scanner scans for the hardware and software inventory information on the Windows inventoried servers. For more information, review the following sections:

- ♦ [“Scanning for the Hardware Inventory Information” on page 535](#)
- ♦ [“Scanning for the Software Inventory Information” on page 537](#)

Scanning for the Hardware Inventory Information

Following are the sources on the inventoried servers from where the hardware inventory information is scanned:

- ♦ [“Desktop Management Interface \(DMI\)” on page 535](#)
- ♦ [“Windows Management Instrumentation \(WMI\)” on page 536](#)
- ♦ [“Probe” on page 536](#)

For more information about the hardware information collected by the Inventory scanner, see [Appendix K, “Hardware Information Collected by the Inventory Scanners,” on page 729](#).

Desktop Management Interface (DMI)

The scanners for scanning the inventoried servers also include scanning based on the industry-standard Desktop Management Interface (DMI) specification 2.0. These programs use the Management Interface (MI) of DMI to look for the hardware components installed on the inventoried server. The scanners scan for specific components that are instrumented on the inventoried server through DMI. The scanners query the DMI service layer to retrieve this information.

The MI allows the DMI-compliant scanners to probe the Service Provider within the Service Layer. The Service Provider collects information from the manageable components and stores the collected information in the Management Information Format database. The Component Interface (CI) communicates with the manageable components and the Service layer. The following figure shows the scanner interaction with DMI.

For more information on DMI standards, see the [DMTF Web site \(http://www.dmtf.org\)](http://www.dmtf.org).

NOTE: If the inventoried servers are DMI compliant and if the Enable DMI Scan check box is selected in the Server Inventory policy, the scanners collect hardware information by querying the DMI Service Layer. Otherwise, the scanners probe for the hardware.

We recommend that you instrument DMI and also install DMI components that are supplied by the vendors.

For example, if you have a Compaq* Family Deskpro* EN Model-SFF6500 running under Windows 98, download the Management Product software - Compaq Insight Management Desktop Agents software for Windows 98 from the Compaq Web site.

For Dell, access the DM/Desktop Management Utilities software from the Dell Web site.

Windows Management Instrumentation (WMI)

The scanners collect hardware information from Windows inventoried servers based on Microsoft Windows Management Instrumentation (WMI) specification.

WMI is the Microsoft implementation of Web-Based Enterprise Management (WBEM) that enables accessing management information in an enterprise environment. WMI 1.5 is fully compliant with Common Information Model (CIM) schema, which is an industry standard. For more information, see [Microsoft WMI Web site \(http://www.microsoft.com/hwdev/driver/WMI\)](http://www.microsoft.com/hwdev/driver/WMI). WMI also works with existing management standards, such as DMI and SNMP.

The scanners use WMI to look for the hardware components installed on the inventoried server. The scanners also scan for specific components that are instrumented on the inventoried server through WMI.

WMI-compliant scanners are supported on Windows inventoried servers only.

You can view the WMI information of the inventoried servers in the Server Inventory.

To obtain WMI information from the inventoried server, you must first download Microsoft's Windows Management Instrumentation - Core Software Installation from [Microsoft WMI Web site \(http://msdn.microsoft.com/downloads/default.asp?url=/downloads/sample.asp?url=/msdn-files/027/001/576/msdncompositedoc.xml\)](http://msdn.microsoft.com/downloads/default.asp?url=/downloads/sample.asp?url=/msdn-files/027/001/576/msdncompositedoc.xml), and then install WMI Core Software on Windows 98 servers

IMPORTANT: Only the WMI Core Software Installation download is required to instrument an inventoried server for WMI. To troubleshoot any WMI related problems, you can use the WMI SDK download. Also, on Windows 2000 servers, the WMI Core Software is already installed.

By default, both DMI and WMI scanning are enabled. To disable either DMI or WMI scanning, deselect the *Enable DMI or Enable WMI* check box in the Inventory policy window.

Probe

Probe is a special built-in algorithm in the Inventory scanner, which is used to collect software information.

Scanning for the Software Inventory Information

The Inventory scanner scans for the following software inventory information on the Windows inventoried servers:

- ◆ “Installed Software Information” on page 537
- ◆ “Disk Usage” on page 538
- ◆ “File Information” on page 538
- ◆ “AntiVirus Definition Files” on page 538

Installed Software Information

The scanner collects software information from the following sources on the inventoried works station:

MSI: Includes software that are installed on the inventoried servers using the Microsoft Installer.

Add-Remove Programs: Includes software that are listed in the Add/Remove Programs window.

Dictionary-based scan: Includes software that is collected based on the software dictionary rules. For more information, see [Section 16.3, “Customizing the Software Inventory Information To Be Scanned For the ZENworks 7 Inventoried Servers,”](#) on page 586

Probe: Probe is a special built-in algorithm in the Inventory scanner. It is used to collect software information about Windows operating system, Internet Explorer, Media Player, Outlook Express, Microsoft Office and its installed components, Novell Client and its installed components, and ZENworks suite and its installed components.

[Table 14-2](#) shows the software information collected by the scanner from the various sources:

Table 14-2 Software Information collected by the Windows Inventory Scanner

| Scanned Attributes | MSI | Add/Remove Programs | Dictionary-based scan | Probe |
|--------------------------|-----|---------------------|-----------------------|-------|
| Product Name | Yes | Yes | Yes | Yes |
| Vendor Name | Yes | No | Yes | Yes |
| Product Version | Yes | Yes | Yes | Yes |
| Product Identifier | Yes | Yes | No | No |
| Product Install Location | Yes | Yes | Yes | Yes |
| Category | No | No | Yes | No |
| Description | No | No | Yes | No |
| Help Link | Yes | Yes | No | No |
| MSI Package GUID | Yes | Yes | No | Yes |
| Display/Internal Version | Yes | Yes | Yes | Yes |
| Language | Yes | Yes | No | Yes |
| UnInstall String | Yes | Yes | No | No |

| Scanned Attributes | MSI | Add/Remove Programs | Dictionary-based scan | Probe |
|---------------------|-----|---------------------|-----------------------|-------|
| Installation Source | Yes | Yes | No | No |
| Display Name | Yes | Yes | No | Yes |
| Support Pack | No | No | No | Yes |
| Product Edition | No | No | No | Yes |
| Last Execution Time | No | Yes | No | No |
| Usage Count | No | Yes | No | No |

Disk Usage

The scanner collects the total disk usage information for the file extensions that are configured in the Software Dictionary editor. For more information, see [Section 16.3, “Customizing the Software Inventory Information To Be Scanned For the ZENworks 7 Inventoried Servers,” on page 586.](#)

File Information

The scanner reports certain information for the files that match with the dictionary entries, and the files that belong to the unidentified software list, which is configured using the software dictionary rules. For more information, see [Section 16.3.10, “Configuring the Software Dictionary Rules,” on page 596](#)

The scanner reports the following file attributes: FileName, FileSize, LastModifiedTime, InternalName, FileVersion, ProductName, ProductVersion, CompanyName, Language, DirectoryPath, and SoftwareDictionaryID.

AntiVirus Definition Files

The scanner collects information about the latest virus definition date and version that is installed on the inventoried servers for the following product versions:

- Symantec* AntiVirus Corporate Edition 8.0
- Symantec AntiVirus Corporate Edition 9.0
- Symantec AntiVirus Corporate Edition 10.0
- Norton AntiVirus* Corporate Edition for Windows 7.0
- Norton AntiVirus Corporate Edition 7.6.1.0000
- Symantec Norton AntiVirus 2000
- Symantec Norton Internet Security 2002
- Symantec Norton AntiVirus 2003 (9.00)
- Symantec Norton AntiVirus 2003 Professional Edition (9.00)
- Symantec Norton AntiVirus 2004 (10.00)
- Symantec Norton Internet Security 2004 (10.00)
- Symantec Norton AntiVirus 2004 Professional (10.00)
- Symantec Norton Internet Security 2004 Professional (10.00)
- Symantec Norton AntiVirus 2005 Professional (11.00)
- Symantec Norton Internet Security 2005 Professional (11.00)
- Network Associates McAfee* VirusScan* 4.0.3 (Windows 9x)
- Network Associates McAfee VirusScan NT 4.0.3a (Windows NT)

Network Associates McAfee NetShield 4.5.0
Network Associates McAfee VirusScan 4.5.0
Network Associates McAfee VirusScan 4.5.1
Network Associates McAfee VirusScan (McAfee Security Center) 8.0
Network Associates McAfee VirusScan ASaP
Network Associates McAfee VirusScan Enterprise 7.1
Network Associates McAfee VirusScan Enterprise 8.0
Central Command Vexira AntiVirus Guard for Windows XP (2000 + NT) 2.10
Central Command Vexira AntiVirus Windows 95/98
Central Command Vexira AntiVirus NT/2000 Servers
Central Command Vexira AntiVirus Server Edition (6.26.xx.xx)
Sophos Anti-Virus - Windows NT/2000/XP/2003
Sophos Anti-Virus - Windows 95/98
Trend Micro PC-cillin 2002 (9.x)
Trend Micro PC-cillin 2003 (10.x)
Trend Micro Internet Security 11.x (PC-cillin)
Trend Micro Internet Security 2005 12.x (PC-cillin)
Trend Micro Server Protect 5.xx
Trend Micro OfficeScan 5.xx - Client for Windows NT/2000/XP
Trend Micro OfficeScan 5.xx - Client for Windows 9x

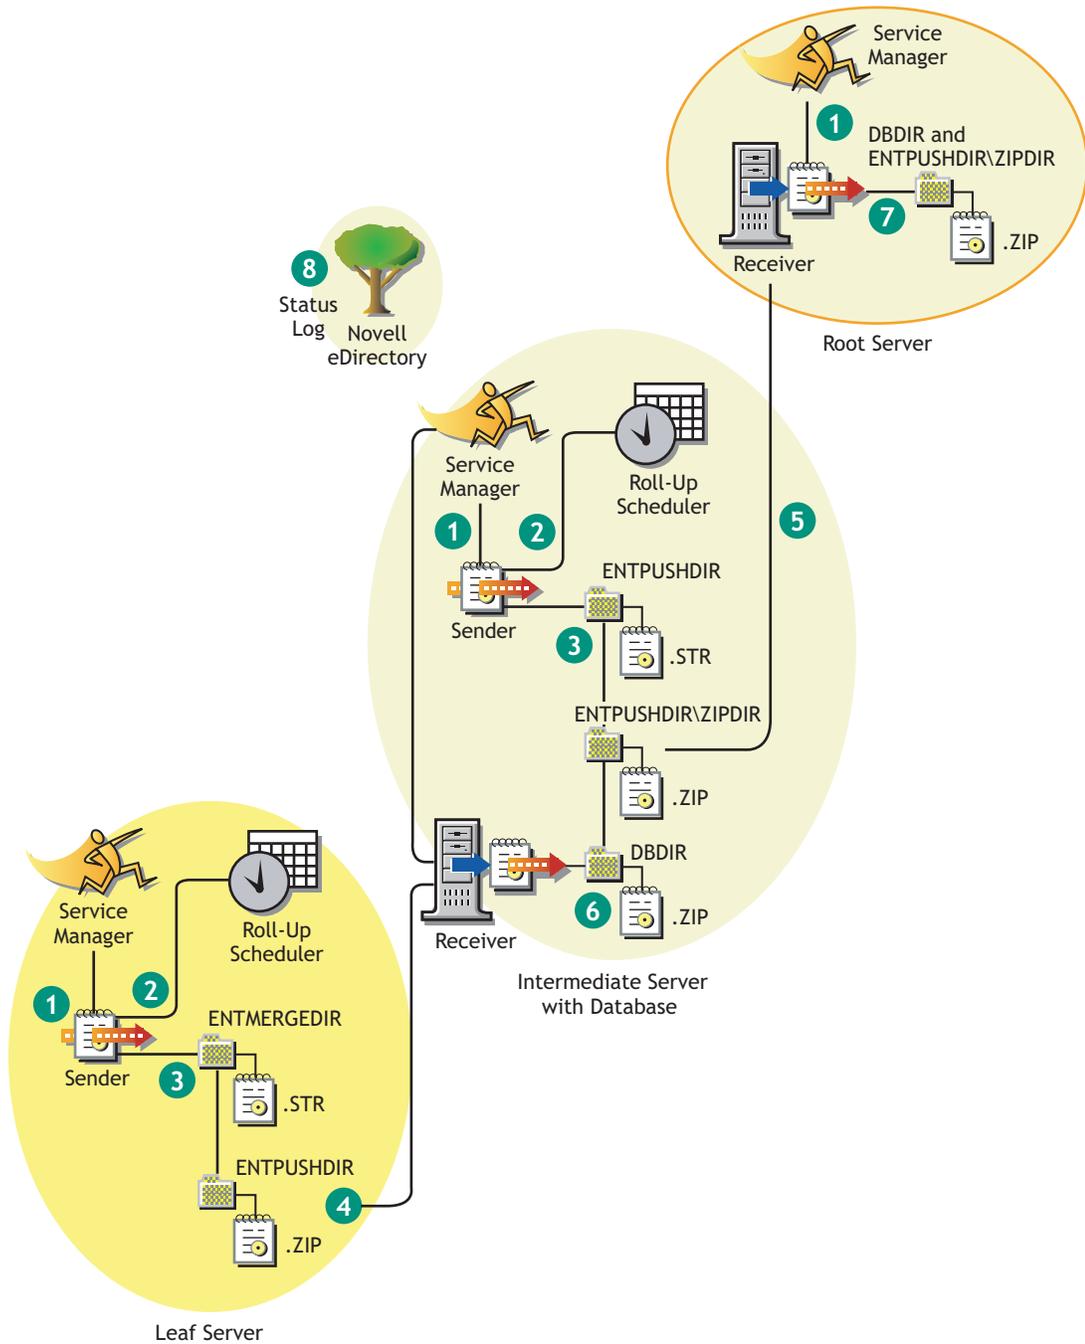
14.4 Understanding the Sender and Receiver

The Sender and the Receiver on the Inventory servers transfer the scan files from the lower-level Inventory servers to the higher-level Inventory servers. The Sender-Receiver uses the ZENworks Web Server to process the XML-RPC requests. The following sections contain more information:

- ◆ [Section 14.4.1, “Understanding the Sender,” on page 541](#)
- ◆ [Section 14.4.2, “Understanding the Receiver,” on page 542](#)
- ◆ [Section 14.4.3, “Understanding the Compressed Scan Data File,” on page 542](#)
- ◆ [Section 14.4.4, “Sender-Receiver Directories,” on page 543](#)

Figure 14-1 depicts the processing done by the Sender-Receiver:

Figure 14-1 Sender-Receiver process



The processing done by the Sender-Receiver is as follows:

1. The Service Manager starts the Sender-Receiver component.
2. The Roll-Up Scheduler activates the Sender at the specified roll-up time.
3. The Sender moves the scan data files (.str) from the enterprise merge directory (entmergedir) to the enterprise push directory (entpushdir) and compresses the files as a .zip file.

4. Each `.zip` file is again compressed with the `.prp` file into a `.zip` file. The `.prp` file is an internal file containing information about the `.zip` file.
5. The Sender sends the `.zip` file from the `\entpushdir` directory to the Receiver on the next-level Inventory server.
6. The Receiver places the `.zip` files to the `\entpushdir\zipdir` directory.
7. The Receiver copies the `.zip` files to the `\entpushdir` directory and deletes the `.zip` files from the `\entpushdir\zipdir` directory.
8. The Receiver copies the `.zip` files to the database directory (`dbdir`) if a database is attached to the Inventory server.
9. The Sender-Receiver logs the status in eDirectory.

14.4.1 Understanding the Sender

The Sender is a Java* component that runs on any Leaf Server or on the Intermediate Server. The Sender is a service loaded by the Service Manager. See [Section 14.9, “An Overview of the Inventory Components on the Inventory Server,” on page 546](#) for a quick reference table of Inventory server components.

The flow of information from the Sender in the roll-up of inventory information is as follows:

1. The Service Manager starts the Sender on the Inventory server. At the specified time scheduled in the Roll-Up Schedule, the Sender moves the scan data files (`.str`) from the enterprise merge directory (`entmergedir`) to the enterprise push directory (`entpushdir`).

The Sender compresses these `.str` files in the `\entpushdir` directory of the Inventory server as a `.zip` file and then deletes the `.str` files. This `.zip` file is again compressed with the `.prp` file into a `.zip` file. The `.prp` file is an internal file containing information about the `.zip` file. For more information, see [“Understanding the Compressed Scan Data File” on page 542](#).

2. The Sender creates a new record in the `zeninvRollUpLog` attribute of the Inventory Service object in eDirectory with the following details: server on which the Sender compresses the `.str` files and the name and size of the `.zip` file.
3. Based on the Discard Scan Data Time in the Inventory Service object properties of the Receiver, the Sender deletes the compressed `.zip` files in the `\entpushdir` directory that have been created earlier than the specified discard scan data time. This removes unwanted scan information being sent in the roll-up.
4. The Sender sends the compressed `.zip` files to the Receiver, with the oldest compressed files sent first.
5. The Sender after transferring the `.zip` file, deletes the compressed files in the `\entpushdir` directory.
6. After the roll-up of information, the Sender updates the `zeninvRollUpLog` attribute of the Inventory server on which the compressed file was created with the following details: Inventory server from which the Sender transmitted the file, name of the `.zip` file, time of transmission, total time taken to transmit the files, and the Inventory server to which it was sent.

In case of rolling up inventory information across trees, the roll-up status messages are logged into the first inventory server receiving the `.zip` file in the tree.

The status information for all actions of the Sender is logged in the Roll-Up Log and Server Status log. For more information, see [“Monitoring Server Inventory Using Status Logs” on page 711](#).

If the Sender is unable to connect to the Receiver, the Sender retries to connect after 10 seconds. The time interval increases exponentially by a factor of 2. After 14 retries, the Sender stops trying to connect to the Receiver. The Sender retries for approximately 23 hours before it discontinues trying. The Sender does not process any other information while it is establishing the connection.

14.4.2 Understanding the Receiver

The Receiver is a Java component that runs on the Intermediate Server or on the Root Server. The Receiver is a service loaded by the Service Manager. See [Section 14.9, “An Overview of the Inventory Components on the Inventory Server,” on page 546](#) for a quick reference table of Inventory server components.

On a Standalone Server, the Receiver is not loaded.

The processing done by the Receiver is as follows:

1. The Receiver receives the scan .zip file from the Sender and places the file in the `\entpushdir\zipdir` directory.
2. The Receiver copies the .zip file to the `\entpushdir` directory and deletes the .zip files from the `\entpushdir\zipdir` directory.

On an Intermediate Server, the file is placed in `\entpushdir`. On an Intermediate Server with Database, or an Intermediate Server with Database and Inventoried Servers, the file is placed in `\entpushdir` and copied to the Database Directory (`dbdir`).

3. The Receiver on the Root Server or the Root Server with Inventoried Servers receives the .zip files from the Senders and places the .zip files in the `\entpushdir\zipdir` directory. It copies the files to the `\dbdir` directory on the Inventory server.
4. The Receiver logs the status information in the Roll-Up log. For more information, see [“Monitoring Server Inventory Using Status Logs” on page 711](#).

14.4.3 Understanding the Compressed Scan Data File

The Sender compresses the scan data files (.str) into a .zip file. This .zip file is again compressed with the .prp file into a .zip file. The .zip file (containing the .zip files and .prp) is named using the following naming conventions:

`scheduledtime_inventoryservername_treename_storedstatus.zip`

where *scheduledtime* refers to the date and time when the .zip file is created, *inventoryservername* refers to the Inventory server on which the .zip file was compressed, *treename* refers to the unique tree name in which the .zip file is currently located, *storedstatus* refers to the storage status of the .zip file, and *ZIP* is the file extension for the compressed files.

The *storedstatus* is represented by 0, 1, or 2. 0 indicates the .zip file has not yet been stored. 1 indicates the .zip file will be stored for the first time in the Inventory server. 2 indicates the .zip file has already been stored once.

The .zip filename changes depending on if the database is attached to the Inventory server.

The .zip file contains the .str files and a property file. The property file is named using the following conventions:

scheduledtime_inventoryservername.prp

The property file contains the scheduled time, Inventory server name, and signature. The signature helps to authenticate the .zip file.

Each .zip file can contain a maximum of 50 .str files.

14.4.4 Sender-Receiver Directories

Table 14-3 provides a quick reference of the directories that the Sender-Receiver uses:

Table 14-3 List of directories used by Sender and Receiver

| Server | Sender | Receiver | Entmergdir | Entpushdir \ Zipdir | Entpushdir | Dbdir |
|--|-------------------------------|-------------------------------|---|---|---|---|
| Leaf Server, Leaf Server with Database | Runs on this Inventory server | -- | Sender moves the .str files to \entpushdir. | -- | Sender compresses the .str files as a .zip file. Sender deletes the .str files. Sends the .zip file to the next-level Inventory server. | -- |
| Intermediate Server | Runs on this Inventory server | Runs on this Inventory server | -- | Receiver receives the .zip files from the lower-level Inventory server in this directory. | Receiver copies the .zip files from the lower-level Inventory server in this directory. Sender sends the .zip files to the next-level Inventory server. | -- |
| Intermediate Server with Inventory Servers | Runs on this Inventory server | Runs on this Inventory server | Sender moves the .str files to \entpushdir. | Receiver receives the .zip files from the lower-level Inventory server in this directory. | Receiver copies the .zip files from \zipdir into this directory. Sender sends the .zip files to the next-level Inventory server. Sender compresses the .str files in to .zip files. Sender deletes the .str files. | -- |
| Intermediate Server with Database | Runs on this Inventory server | Runs on this Inventory server | -- | Receiver receives the .zip files from the lower-level Inventory server in this directory. | Receiver copies the .zip files from \zipdir into this directory. Sender sends the .zip file to the next-level Inventory server. | Receiver copies the file in this directory. |

| Server | Sender | Receiver | Entmergdir | Entpushdir \ Zipdir | Entpushdir | Dbdir |
|---|-------------------------------|-------------------------------|---|---|--|---|
| Intermediate Server with Database and Inventoried Servers | Runs on this Inventory server | Runs on this Inventory server | Sender moves the .str files to \entpushdir. | Receiver receives the .zip files from the lower-level Inventory server in this directory. | Receiver copies the .zip files from \zipdir into this directory. Sender compresses the .str files as a .zip file. Sender deletes the .str files. Sender sends the .zip file to the next-level Inventory server. | Receiver copies the file in this directory. |
| Root Server with Inventoried Servers | -- | Runs on this Inventory server | -- | Receiver receives the .zip files from the lower-level Inventory server in this directory. | -- | Receiver copies the .zip files from the lower-level Inventory server in this directory. |

14.5 Understanding the Selector

The Selector is a Java component on the Inventory server that receives the inventory information from the Inventory servers. These Inventory servers can be any of the following: Leaf Server, Leaf Server with Database, Intermediate Server with Database and Inventoried Servers, Intermediate Server with Inventoried Servers, Root Server with Inventoried Servers, and Standalone Server. See [Section 14.9, “An Overview of the Inventory Components on the Inventory Server,” on page 546](#) for a quick reference table of Inventory server components.

The processing done by the Selector is as follows:

1. While scanning the inventoried server, the Scanner creates a scan data file (.str) in the scan directory (scandir) at the Inventory server for each scan done on the inventoried server. The location of scandir is obtained from the Inventory Service object. The Selector processes the .str files placed by the Scan Collector in the scandir directory.
2. The Selector checks for the following conditions to ensure that the .str file generated by the Scanner is valid:
 - ♦ The integer value that is generated by using the .str file and logged into the .str file by the Scanner and the integer value generated by using the .str file by the Selector should be the same.
 - ♦ The actual size of the .str file should be in sync with the size recorded in the .str file.

The Selector processes only valid .str files. If invalid files are present in the directory, the Selector deletes the files.

- Based on the role of the Inventory server, the Selector copies the `.str` files to the `dbdir` directory (if the database is attached) and the `entmerge` directory. If the `.str` file already exists in the directory, it overwrites the file.

The following table lists the directories that the Selector copies or renames the files to:

| Server | Copies the <code>.str</code> file to the Database Directory (<code>dbdir</code>) | Renames the <code>.str</code> file in the Database Directory (<code>dbdir</code>) | Renames the <code>.str</code> file in the Enterprise Merge Directory (<code>entmergedir</code>) |
|---|--|---|---|
| Leaf Server with Database | Yes | -- | Yes |
| Leaf Server | -- | -- | Yes |
| Intermediate Server with Database and Inventoried Servers | Yes | -- | Yes |
| Standalone Server | -- | Yes | -- |
| Root Server with Inventoried Servers | -- | Yes | -- |

- The Selector logs the status in the Server log. For more information, see [“Monitoring Server Inventory Using Status Logs” on page 711](#).

14.6 Understanding the Storer

The Storer is a Java component on the Inventory server that has a database attached to it. These Inventory servers can be any of the following: Leaf Server with Database, Intermediate Server with Database, Intermediate Server with Database and Inventoried Servers, Root Server, and Root Server with Inventoried Servers. See [Section 14.9, “An Overview of the Inventory Components on the Inventory Server,” on page 546](#) for a quick reference table of Inventory server components.

The Storer runs as a service loaded by the Service Manager. It processes the files in the `\dbdir` directory.

The processing done by the Storer is as follows:

- The Storer reads the Startup configuration parameters from the Inventory server Configuration Service.
- From the Inventory server configuration information stored in eDirectory, the Storer looks in the database directory (`dbdir`) for the scan files. The Inventory server configuration information determines the location of `\dbdir` and the database server from the eDirectory policy. The Selector places the `.str` files in `\dbdir` and the Receiver places the `.zip` files in `\dbdir`.
- The Storer processes the `.str` files and the `.zip` files alternately.
- The Storer extracts the `.zip` file containing the compressed `.str` files and the `.prp` file to a temp directory (`\dbdir\temp`) and updates the database with the inventory information of the `.str` files for the inventoried servers.
- The Storer updates the status in the Inventory server Status log and updates the Roll-Up log. You can view the Inventory server status information in the Inventory server Status log.

In case of rolling up inventory information across trees, the roll-up status messages are logged into the first inventory server receiving the `.zip` file in the tree. For more information, see [“Monitoring Server Inventory Using Status Logs” on page 711](#).

14.7 Understanding the Dictionary Provider and Dictionary Consumer

All Inventory servers and inventoried servers have Dictionary Provider and Dictionary Consumer services that are automatically installed during the Server Inventory installation.

When an inventoried server is started, the Dictionary Consumer reads the Dictionary Update policy and contacts the Dictionary Provider (running on the Inventory server) specified in the policy.

Subsequently, the Dictionary Consumer checks for the dictionary updates based on the schedule set in the Dictionary Update policy. It compares the date of the dictionary file on the Dictionary Provider with the file that has been locally stored. If the file on the Dictionary Provider is newer, then the Dictionary Consumer downloads the file from the Dictionary Provider using XML-RPC as per the schedule.

14.8 Understanding the Upgrade Service

The Upgrade service runs as a service loaded by the Service Manager and performs the following functions:

1. Migrates ZENworks for Servers 3.0.2 database to ZENworks 7 Server Management database.
2. Converts the ZENworks for Servers 3.0.2 residue `.str` files to ZENworks 7 SP1 Server Management with SP1 `.str` files.

The Upgrade service corrects the Inventory database schema and information to make it compatible with ZENworks 7 Server Management with SP1 and ZENworks 7 Desktop Management. The Upgrade service performs all the functions in a state-driven method. This is to make sure that the Upgrade service does not execute the same steps when one step is executed successfully. The Upgrade service runs as an uninterrupted service. Therefore, you cannot manually stop the Upgrade service. The Upgrade service stops automatically after completing all its functions.

The Database migration activity is additionally traced into a migration log, which could be found in the `installation_path\zenworks\inv\server\wminv\logs\migrationlogs` directory.

14.9 An Overview of the Inventory Components on the Inventory Server

Depending on the type of the Inventory server, the inventory components exist on the Inventory server as listed in [Table 14-4](#).

Table 14-4 Inventory components running on the Inventory server

| Server Component | Root Server | Root Server with Inventoried Servers | Leaf Server | Leaf Server with Database | Intermediate Server | Intermediate Server with Database and Inventoried Servers | Intermediate Server with Database | Intermediate Server with Inventoried Servers | Standalone Server |
|---|-------------|--------------------------------------|-------------|---------------------------|---------------------|---|-----------------------------------|--|-------------------|
| Service Manager | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Scan Collector | No | Yes | Yes | Yes | No | Yes | No | Yes | Yes |
| Selector | No | Yes | Yes | Yes | No | Yes | No | Yes | Yes |
| Storer | Yes | Yes | No | Yes | No | Yes | Yes | No | Yes |
| Sender | No | No | Yes | Yes | Yes | Yes | Yes | Yes | No |
| Receiver | Yes | Yes | No | No | Yes | Yes | Yes | Yes | No |
| Database | Yes | Yes | No | Yes | No | Yes | Yes | No | Yes |
| Inventory Removal Service | Yes | Yes | Yes | Yes | No | Yes | Yes | Yes | Yes |
| Upgrade Service | Yes | Yes | No | Yes | No | Yes | Yes | No | Yes |
| Dictionary Consumer and Dictionary Provider | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |

14.10 Understanding the Inventory Database

Server Inventory provides a centralized Common Information Model (CIM)-compliant Sybase database. The Inventory database serves as a repository of hardware and software information for the servers. The network administrator can view the inventory information, query the database, and generate inventory reports in ConsoleOne. For more information, see [Chapter 15, “Understanding the ZENworks 7 Server Managements Inventory Database Schema,”](#) on page 549

Understanding the ZENworks 7 Server Managements Inventory Database Schema

15

This section describes the design of the Novell® ZENworks® 7 Server Management Inventory database schema implemented using the Common Information Model (CIM) of the Distributed Management Task Force (DMTF). To understand this section effectively, you should be familiar with terminology such as CIM and Desktop Management Interface (DMI). You should also have a solid understanding of Relational Database Based Managed Systems (RDBMS) and database concepts.

The following sections provide in-depth information:

- ♦ [Section 15.1, “Overview,” on page 549](#)
- ♦ [Section 15.2, “CIM Schema,” on page 550](#)
- ♦ [Section 15.3, “Inventory Database Schema in ZENworks 7 Server Management,” on page 556](#)

15.1 Overview

The DMTF is the industry organization leading the development, adoption, and unification of management standards and initiatives for desktop, enterprise, and Internet environments. For more information about DMTF, see the [DMTF Web site \(http://www.dmtf.org\)](http://www.dmtf.org).

The DMTF CIM is an approach to system and network management that applies the basic structuring and conceptualization techniques of the object-oriented paradigm. The approach uses a uniform modeling formalism that together with the basic repertoire of object-oriented constructs supports the cooperative development of an object-oriented schema across multiple organizations.

A management schema is provided to establish a common conceptual framework at the level of a fundamental topology, both with respect to classification and association, and to a basic set of classes intended to establish a common framework for a description of the managed environment. The management schema is divided into the following conceptual layers:

- ♦ **Core Model:** An information model that captures notions that are applicable to all areas of management.
- ♦ **Common Model:** An information model that captures notions that are common to particular management areas, but independent of a particular technology or implementation. The common areas are systems, applications, databases, networks, and devices. The information model is specific enough to provide a basis for the development of management applications. This model provides a set of base classes for extension into the area of technology-specific schema. The Core and Common models together are expressed as the CIM schema.
- ♦ **Extension Schema:** This schema represents technology-specific extensions of the Common model. These schema are specific to environments, such as operating systems, for example, NetWare® or Microsoft Windows.

CIM comprises a specification and a schema (see the [DMTF Web site \(http://www.dmtf.org/standards/standard_cim.php\)](http://www.dmtf.org/standards/standard_cim.php)). The specification defines the meta-schema plus a concrete representation language called Managed Object Format (MOF).

15.2 CIM Schema

The elements of the meta schema are classes, properties, and methods. The meta schema also supports indications and associations as types of classes and references as types of properties.

Classes can be arranged in a generalization hierarchy that represents subtype relationships between classes. The generalization hierarchy is a rooted, directed graph that does not support multiple inheritance.

A regular class may contain scalar or array properties of any intrinsic type such as Boolean, integer, string, and others. It cannot contain embedded classes or references to other classes.

An association is a special class that contains two or more references. It represents a relationship between two or more objects. Because of the way associations are defined, it is possible to establish a relationship between classes without affecting any of the related classes. That is, addition of an association does not affect the interface of the related classes. Only associations can have references.

The schema fragment in [Figure 16-3](#) shows the relationships between some CIM objects that ZENworks 7 Server Management uses:

Figure 15-1 CIM Objects as Used in ZENworks 7 Server Management

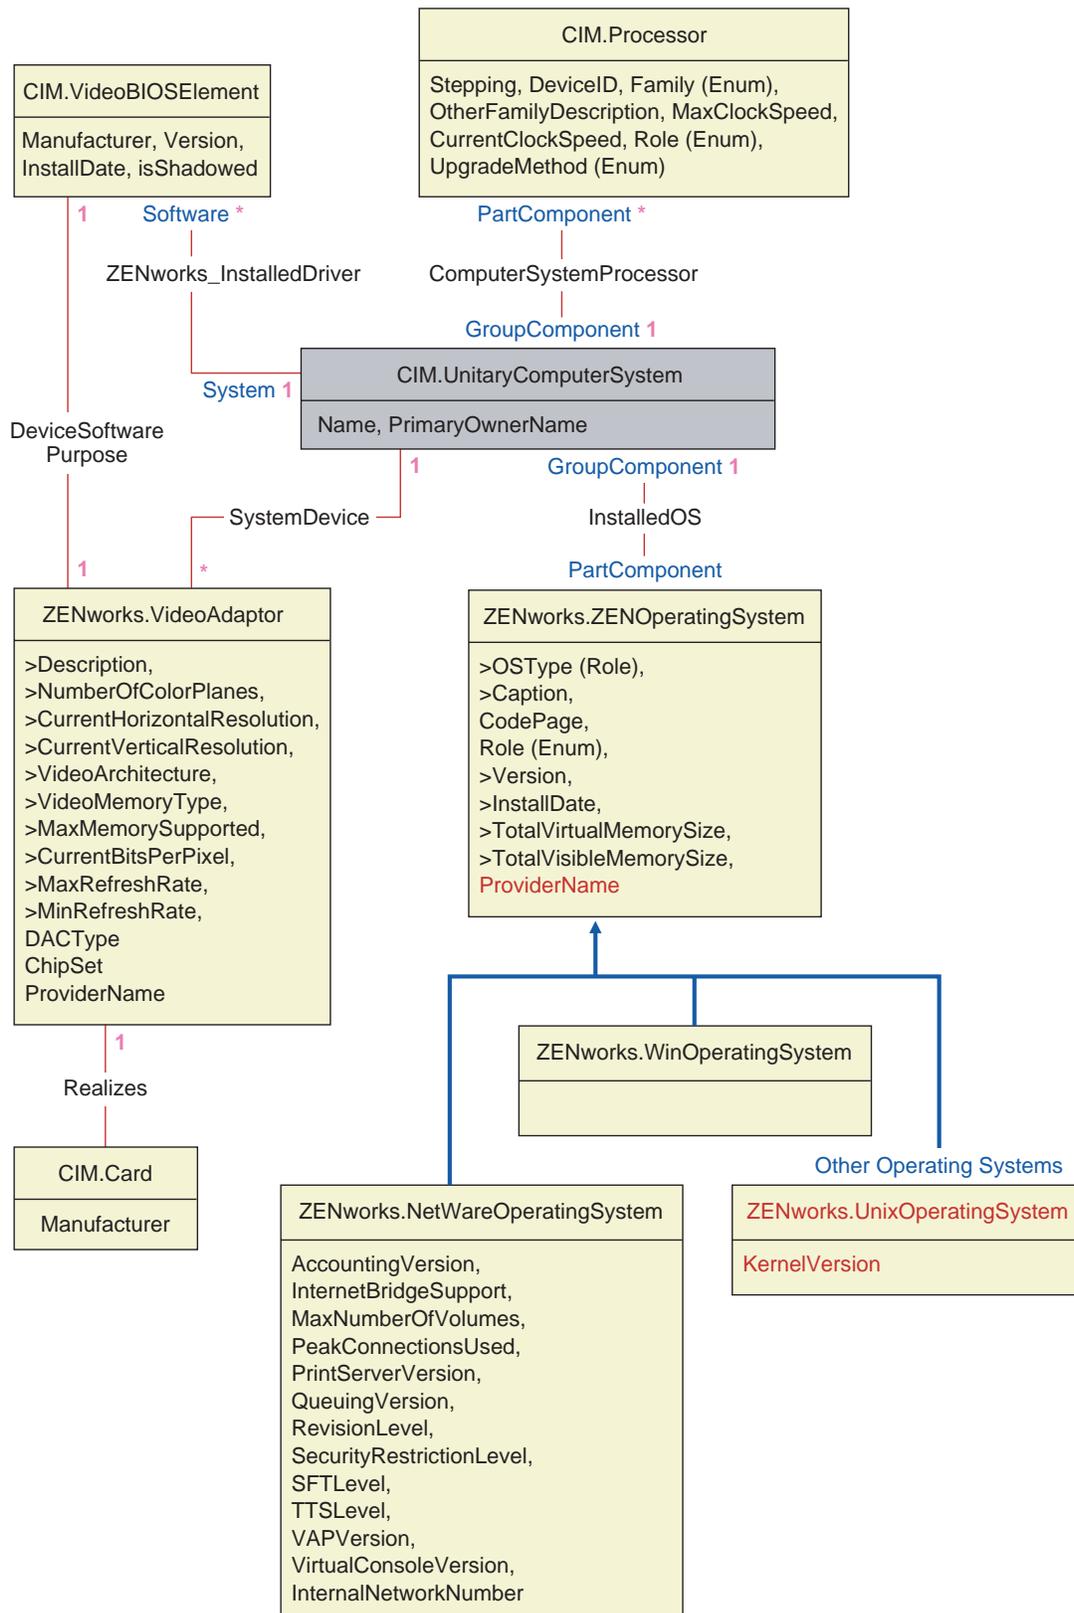
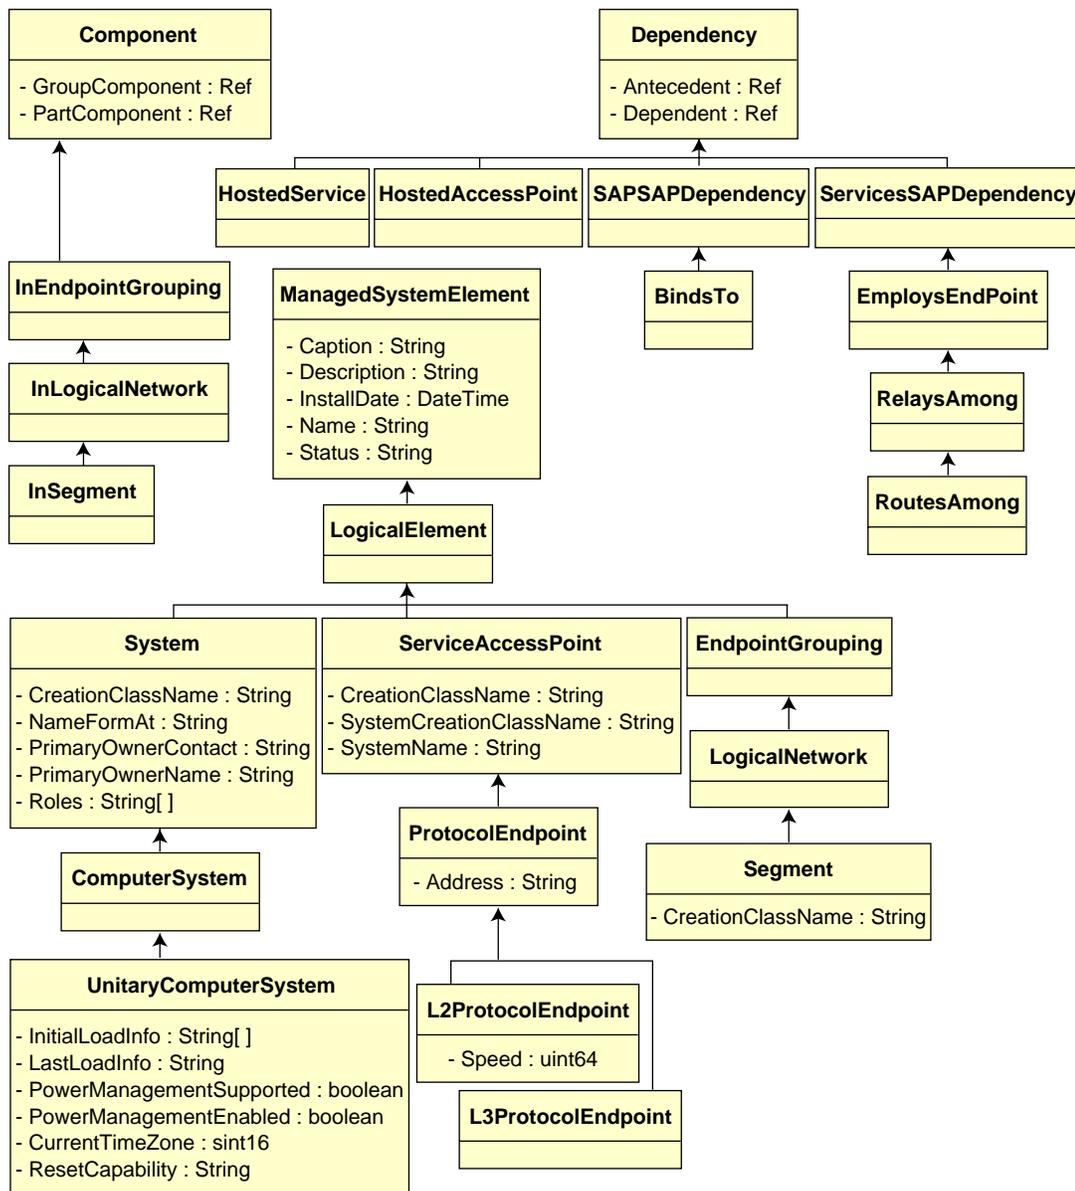


Figure 15-1 shows how the CIM schema maps to a relational DBMS schema. The classes are shown with the class name as the box heading. The associations are labeled within the lines between two classes.

The inheritance hierarchy of this schema fragment is shown in Figure 15-2 of the CIM 2.2 schema. The references shown as type Ref are in bold with each association sub-type narrowing the type of the reference.

Figure 15-2 Inheritance Hierarchy



15.2.1 CIM-to-Relational Mapping

CIM is an object model complete with classes, inheritance, and polymorphism. The generated mapping to a relational schema preserves these features to the maximum extent. The following two aspects are part of the relational mapping:

- ♦ **Logical Schema:** The logical schema defines how the information appears to applications, similar to an API. The goal is that the logical schema remains the same irrespective of the underlying database so that application software can run unchanged on any supported databases. Though SQL is a standard, this goal is not fully possible. Application software will need to know more about the database in use and this information can be abstracted and isolated to a small area of the application code.
- ♦ **Physical Schema:** The physical schema defines how the information is structured in the database. The schema tends to be specific to the database because of the nature of SQL and RDBMS. This document will describe the physical schema in general terms only.

A table in the database represents each class in the CIM hierarchy. A column of the appropriate type in the table represents each non-inherited property in the class. Each table also has a primary key, `id$`, which is a 64-bit integer that uniquely identifies an instance. An instance of a CIM class is represented by a row in each table that corresponds to a class in its inheritance hierarchy. Each row has the same value for `id$`.

Each CIM class is also represented by a view that uses `id$` to join rows from the various tables in the inheritance hierarchy to yield a composite set of properties (inherited plus local) for an instance of that class. The view also contains an extra column, `class$`, of type integer that represents the type of the actual (leaf-most) class of the instance.

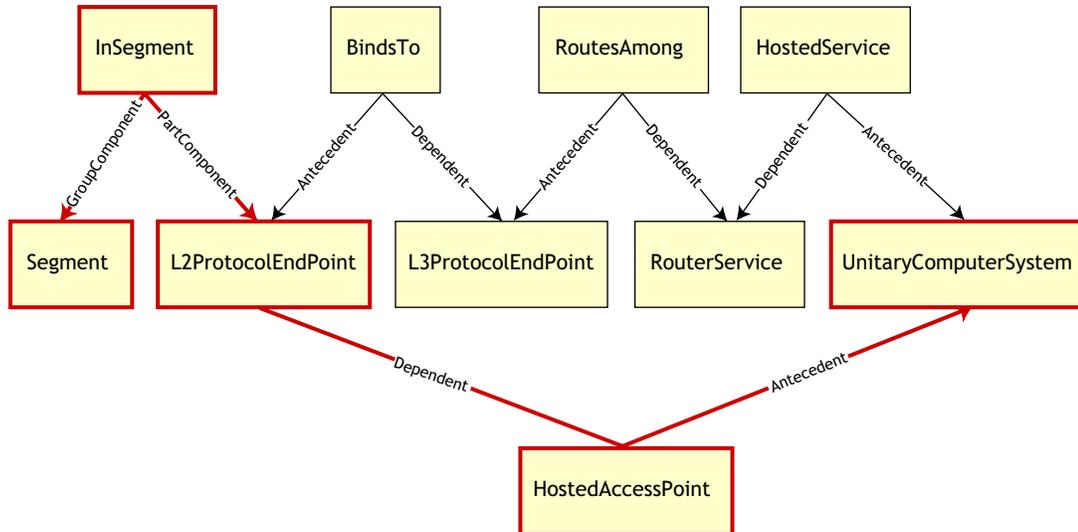
Associations are mapped in the same manner as regular classes, with a reference property being represented by a column with the `id$` field of the referenced object instance. Thus, associations can be traversed by doing a join between the reference field in the association and the `id$` field in the referenced table.

Figure 15-3 depicts a typical query using this mapping:

Figure 15-3 Query Using CIM-to-Relational Mapping

Get Computers for Segment

```
SELECT CIM.UnitaryComputerSystem.*
FROM CIM.UnitaryComputerSystem, CIM.Segment, CIM.L2ProtocolEndPoint,
      CIM.HostedAccessPoint, CIM.InSegment
WHERE CIM.SegmentName = 'xxx'
AND CIM.InSegment.GroupComponent = CIM.Segment.id$
AND CIM.InSegment.PartComponent = CIM.L2ProtocolEndPoint.id$
AND CIM.HostedAccessPoint.Dependent = CIM.L2ProtocolEndPoint.id$
AND CIM.HostedAccessPoint.Antecedent = CIM.UnitaryComputerSystem.id$
```



This query finds all the computers attached to a given network segment. The classes and relationships involved are highlighted with borders.

The following topics describe both the schema types:

- ◆ “Logical Schema” on page 554
- ◆ “Physical Schema” on page 556

Logical Schema

The logical schema is the database schema as seen by users of the database and the application program. The schema consists of stored procedures and views. The underlying tables are not visible to the application.

Inventory components use JDBC to issue SQL statements to the RDBMS and to convert between RDBMS data types and Java* data types. The use of JDBC with stored procedures and views provides a level of abstraction that insulates application code from the underlying database technology and from changes to the physical schema.

The various elements of the logical schema are discussed in more detail in the following sections:

- ◆ “Naming Schema Elements” on page 555
- ◆ “Users and Roles” on page 555
- ◆ “Data Types” on page 556

- ◆ “Views” on page 556

Naming Schema Elements

We recommend that you use the CIM names unchanged in the database schema. Some problems may still ensue because of the differences in the naming schemes, such as the following:

- ◆ Names in CIM and SQL are not case sensitive.
- ◆ All databases have different sets of reserved words that must be enclosed in quotes (" ") when used as schema element names; however, in Oracle, enclosing a name in quotes makes it case sensitive.
- ◆ CIM classes avoid using SQL reserved words as names.
- ◆ CIM names are not limited in length and usually the names are long. Sybase allows up to 128 characters, but Oracle restricts the names to 30 characters.

Most of these problems are avoided during schema generation by preserving the case of CIM names, abbreviating any names longer than 30 characters, and placing quotes around any name that is in the union of the sets of reserved words.

Any name longer than 28 characters is abbreviated to a root name of 28 or fewer characters to allow a two-character prefix so that all associated SQL schema elements can use the same root name. The abbreviation algorithm shortens a name so that it is mnemonic, recognizable, and also unique within its scope. The abbreviated name is given a # character as a suffix (note that # is an illegal character in CIM) to prevent clashes with other names. If two or more names within the same scope generate the same abbreviation, an additional digit is appended to make the name unique. For example, AttributeCachingForRegularFilesMin is abbreviated to AttCacForRegularFilesMin#.

All such mangled names are written to the mangled name table so that a program can look up the real CIM name and retrieve the mangled name to use with the SQL.

Views are the schema elements that are most often manipulated by application code and queries. They use the same name as the CIM class they represent. For example, the CIM.UnitaryComputerSystem class is represented by a view named CIM.UnitaryComputerSystem.

When necessary, names for indexes and auxiliary tables are created by concatenating the class name and property name separated by a \$ character. These names are usually abbreviated. For example, NetworkAdapter\$NetworkAddresses is abbreviated to NetAdapter\$NetAddresses#. This does not have any adverse impact on ZENworks 7 Server Management schema users.

Users and Roles

In SQL, a user with the same name as the schema is the owner of each schema, for example, CIM, ManageWise[®], ZENworks[®], and others.

Additionally, there is an MW_DBA user that has Database Administrator privileges and rights to all schema objects. The MW_Reader role has read-only access to all schema objects and the MW_Updater role has read-write-execute access to all schema objects.

Application programs should access the database as either MW_Reader or MW_Updater for a Sybase database, MWO_Reader or MWO_Updater for an Oracle database, and MWM_Reader or MWM_Updater for MS SQL Server database depending on their requirements.

Data Types

CIM data types are mapped to the most appropriate data type provided by the database. Usually, the Java application does not require the type because it uses JDBC to access the data.

Java does not natively support unsigned types, so you should use classes or integer types of the next size to represent them. Also, ensure that there are no problems while reading or writing to the database. For example, reading or writing a negative number to an unsigned field in the database is likely cause an error.

Strings in CIM and Java are Unicode*, so the database is created using the UTF-8 character set. Internationalization does not pose any problems; however, it may create problem with case sensitivity in queries.

All databases preserve the case of string data stored within them, but may access the data as either case sensitive or otherwise during queries. In ZENworks 7 Server Management, the Inventory Query component is not affected because the queried data is retrieved from the database before being queried and so case sensitivity is automatically taken care of.

In CIM, strings may be specified with or without a maximum size in characters. Many strings have no specified size, which means they can be unlimited in size. For efficiency reasons, these unlimited strings are mapped to a variable string with maximum size of 254 characters. CIM strings with a maximum size are mapped to variable database strings of the same size. The size in the database is in bytes and not as characters because a Unicode character may require more than one byte for storage.

Views

Each CIM class is represented in the database by a view that contains all the local and inherited non-array properties of that class. The view is named the same as the CIM class.

Views can be queried using the SELECT statement and updated using the UPDATE statement. Because views cannot be used with the INSERT and DELETE statements, use the constructor and destructor procedures.

Physical Schema

- ◆ Table definitions 't\$xxx'
- ◆ Index definitions 'i\$xxx'
- ◆ Trigger definitions 'x\$xxx', 'n\$xxx' and 'u\$xxx'
- ◆ Sequence definitions (Oracle) 's\$xxx'
- ◆ Stored procedures and functions

The logical schema is layered on top of the physical schema and makes it unnecessary for users and applications to know the physical schema.

15.3 Inventory Database Schema in ZENworks 7 Server Management

The following section describes the database schema classes and the extensions and associations made to the CIM schema for use in ZENworks 7 Server Management. These extensions have ZENworks or ManageWise as their schema name. ZENworks.*classname* refers to the extended class

in the ZENworks schema and `ManageWise.classname` refers to the extended class in the `ManageWise` schema.

The following sections help you understand the ZENworks 7 Server Management database schema:

- ♦ [Section 15.3.1, “Case Study of CIM Schema Implementation in ZENworks 7 Server Management,” on page 557](#)
- ♦ [Section 15.3.2, “Legends for Schema Diagrams,” on page 559](#)
- ♦ [Section 15.3.3, “Schema Diagrams of CIM and the Extension Schema in ZENworks 7 Server Management,” on page 559](#)
- ♦ [Section 15.3.4, “Software Inventory Schema,” on page 567](#)
- ♦ [Section 15.3.5, “Sample Inventory Database Queries,” on page 573](#)

15.3.1 Case Study of CIM Schema Implementation in ZENworks 7 Server Management

The following scenario describes an inventoried server that has two parallel ports with a specified interrupt number.

In the following schema diagram, the `CIM.UnitaryComputerSystem` represents a managed inventory system.

In this illustration, class `CIM.PointingDevice` associates to `CIM.UnitaryComputerSystem` using the association `CIM.SystemDevice` with `SystemDevice.GroupComponent` pointing to `CIM.UnitaryComputerSystem` and `SystemDevice.PartComponent` pointing to `CIM.PointingDevice`. The relationship between the two classes is one to many. This means a computer system might have more than one pointing devices.

Class `CIM.IRQ` associates to `CIM.PointingDevice` using the association `CIM.AllocatedResource`. `Dependent` pointing to `CIM.PointingDevice` and `Antecedent` pointing to `CIM.IRQ`.

Class `ZENworks.ZENKeyboard` associates to `CIM.UnitaryComputerSystem` using the association `CIM.SystemDevice` with `SystemDevice.GroupComponent` pointing to `CIM.UnitaryComputerSystem` and `SystemDevice.PartComponent` pointing to `ZENworks.ZENKeyboard`. The relationship between the two classes is one to one. This means a computer system can have only one Keyboard.

Class `ZENworks.BIOS` associates to `CIM.UnitaryComputerSystem` using the association `CIM.SystemBIOS` with `SystemDevice.GroupComponent` pointing to `CIM.UnitaryComputerSystem` and `SystemBIOS.PartComponent` pointing to `ZENworks.BIOS`. The relationship between the two classes is one to one. This means a computer system can have only one BIOS.

Class `CIM.ZENworks.ParallelPort` associates to `CIM.UnitaryComputerSystem` using the association `CIM.SystemDevice` with `SystemDevice.GroupComponent` pointing to `CIM.UnitaryComputerSystem` and `SystemDevice.PartComponent` pointing to `CIM.ZENworks.ParallelPort`. The relationship between the two classes is one to many. This means a computer system might have more than one parallel port.

Class `ZENworks.BUS` associates to `CIM.UnitaryComputerSystem` using the association `CIM.SystemDevice` with `SystemDevice.GroupComponent` pointing to `CIM.UnitaryComputerSystem` and `SystemDevice.PartComponent` pointing to `ZENworks.BUS`. The

relationship between the two classes is one to many. This means a computer system can have more than one bus.

Class ManageWise.User associates to CIM.UnitaryComputerSystem using CurrentLoginUser and LastLoginUser. In the CurrentLoginUser association, the specific instance of User is the one who is currently logged into the inventoried server. In the LastLoginUser association, the specific instance of User is the one who logged last into the inventoried server.

Class CIM.IRQ associates to CIM.ParallelPort using the association CIM.AllocatedResource. Dependent pointing to CIM.ParallelPort and Antecedent pointing to CIM.IRQ.

Figure 15-4 CIM Schema Implementation

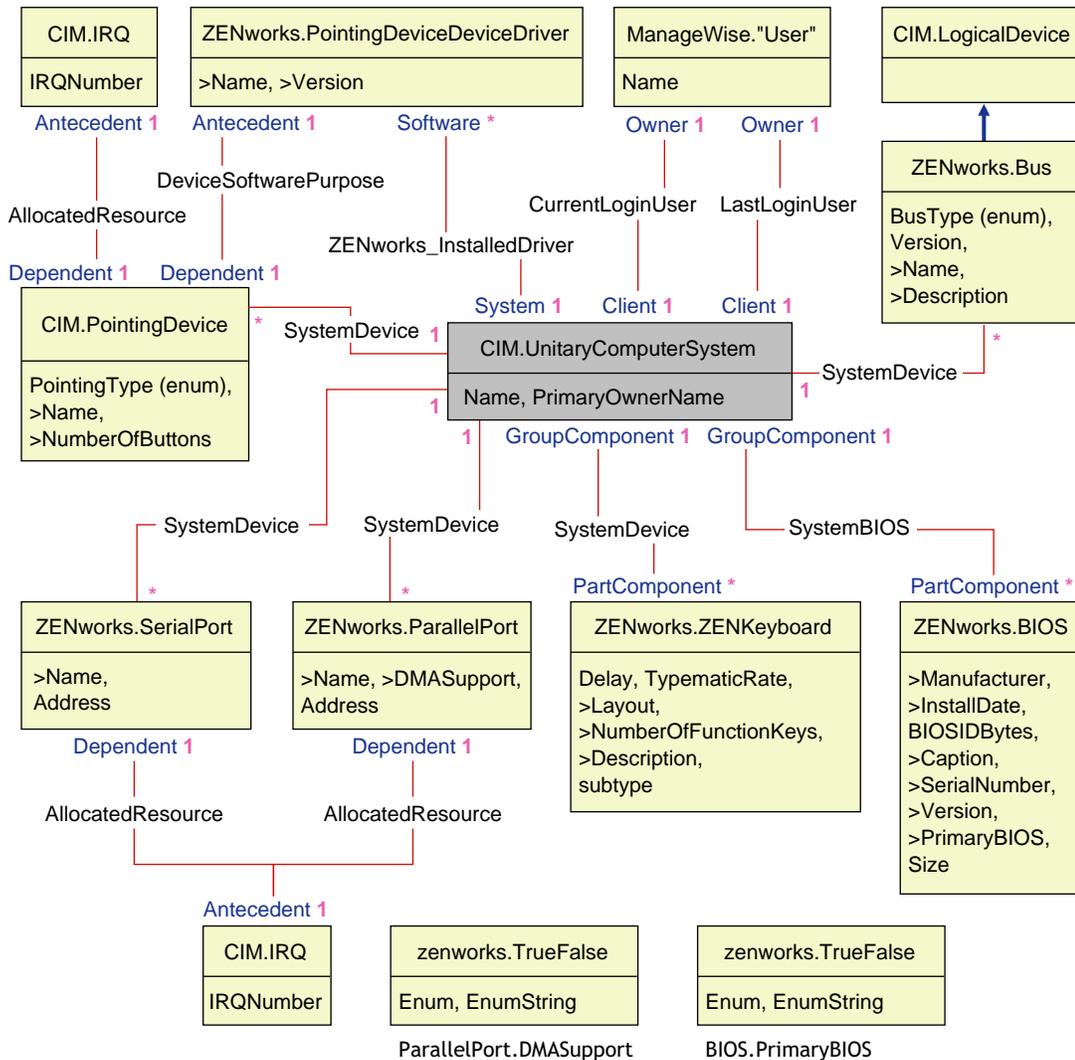


Figure 15-4 illustrates the following:

- All components that a computer system manages are represented as associations from the UnitaryComputerSystem class. The type of references (1..n, 1..1) between two classes are marked.
- Those associations that do not have a schema name are assumed as CIM schema.

There are three instances of ZENworks.ParallelPort associated to one instance of CIM.UnitaryComputerSystem using three instances of CIM.SystemDevice associations. CIM.SystemDevice.GroupComponent references UnitaryComputerSystem and CIM.SystemDevice.PartComponent references ParallelPort.

This is called 1 to n object reference relationship and is depicted in the illustration as 1..*. Similarly, every instance of ParallelPort has a corresponding instance of CIM.IRQ designating the port's irq. This is one-to-one relationship and is depicted as 1..1.

All other classes follow similar representation.

For schema diagrams of other classes, see “[Schema Diagrams of CIM and the Extension Schema in ZENworks 7 Server Management](#)” on page 559.

15.3.2 Legends for Schema Diagrams

The legends for reading the schema diagrams are as follows:

- ◆ Class names are enclosed in boxes with the class name as the heading and the attribute names within it.
- ◆ Red lines connect two classes using an association class.
- ◆ Blue lines indicate the class inheritance hierarchy. The class pointed by the arrow is the class that is being inherited from. The class from where the arrow emanates is the inheriting class.
- ◆ The association class name is shown within the line joining two classes.
- ◆ References of the association class are marked on either side of the associated classes.

For an explanation about CIM schema, see the CIM 2.2 schema specification on the [DMTF Web site \(http://www.dmtf.org\)](http://www.dmtf.org).

15.3.3 Schema Diagrams of CIM and the Extension Schema in ZENworks 7 Server Management

The schema diagrams of the CIM and extension schema on the following pages model the Inventory database in ZENworks 7 Server Management.

Figure 15-5 Schema for Processor, Operating Systems, and Video Adapter

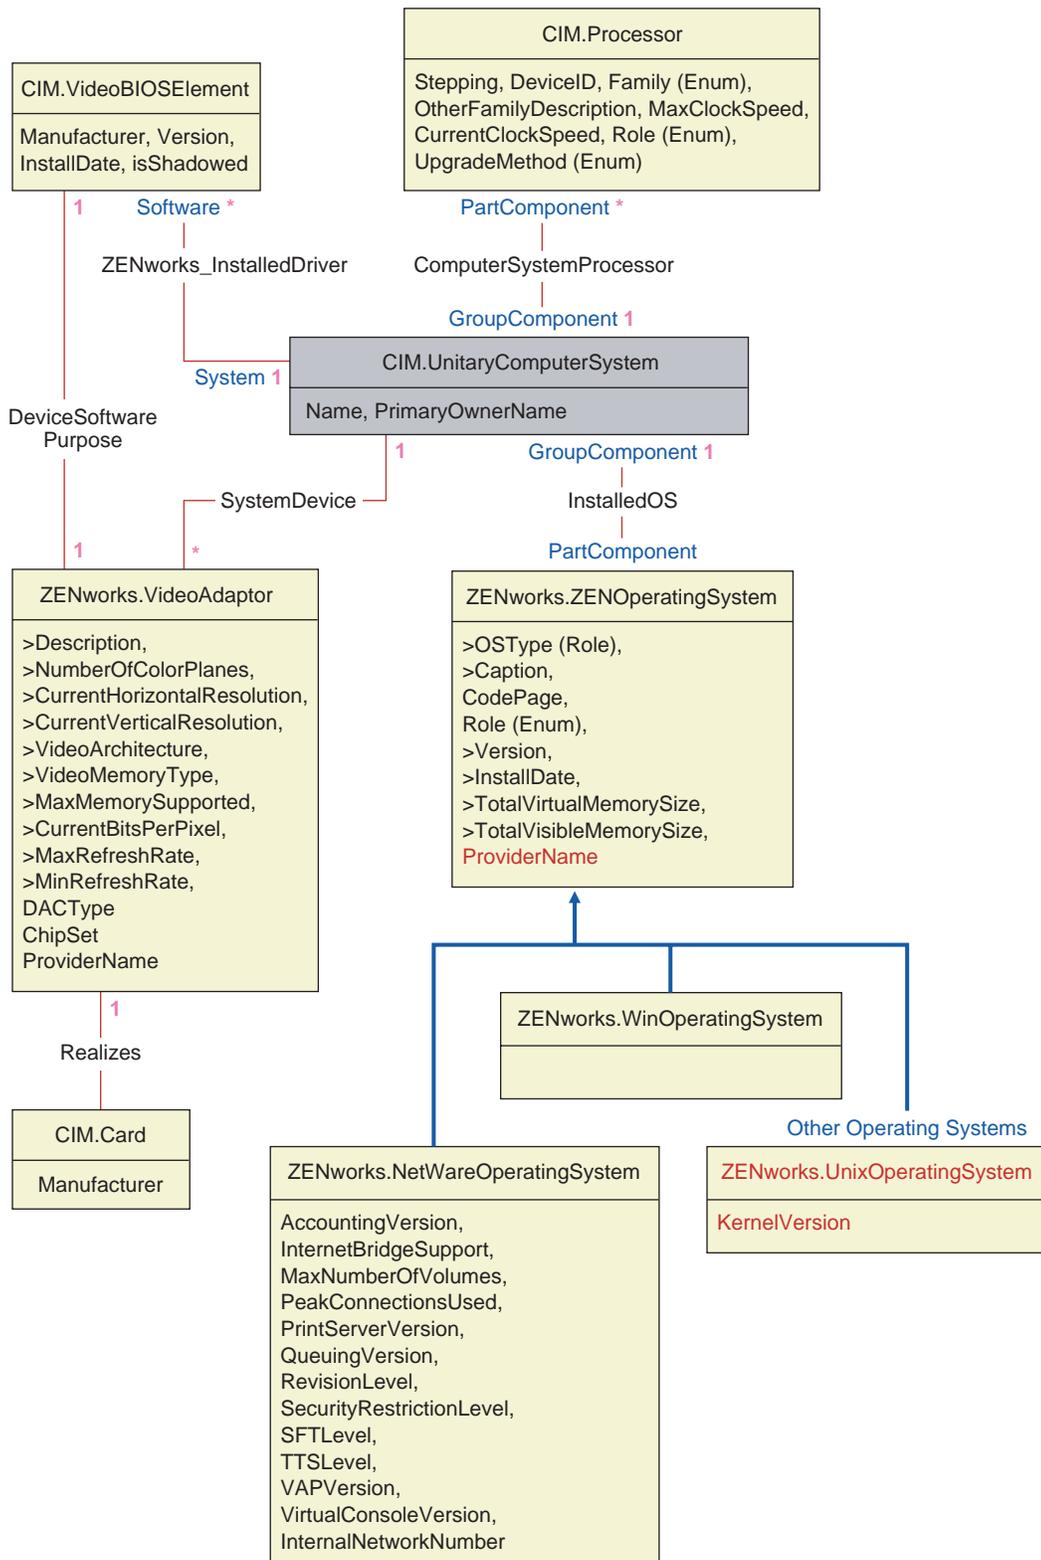


Figure 15-6 Schema for Inventory Scanner and NetWare Client

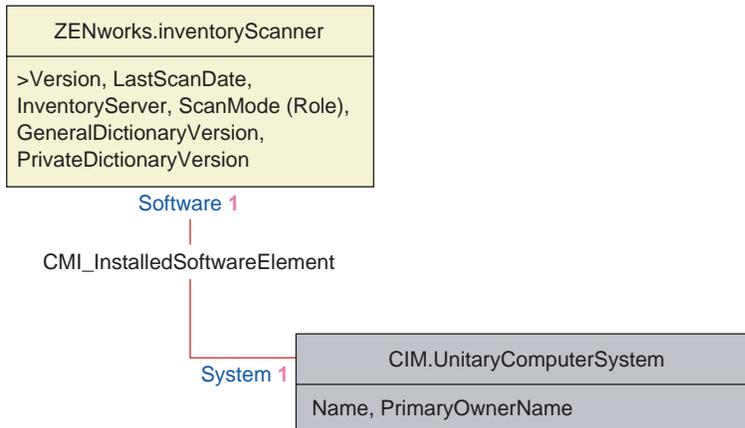


Figure 15-7 Schema for Chassis and System Information

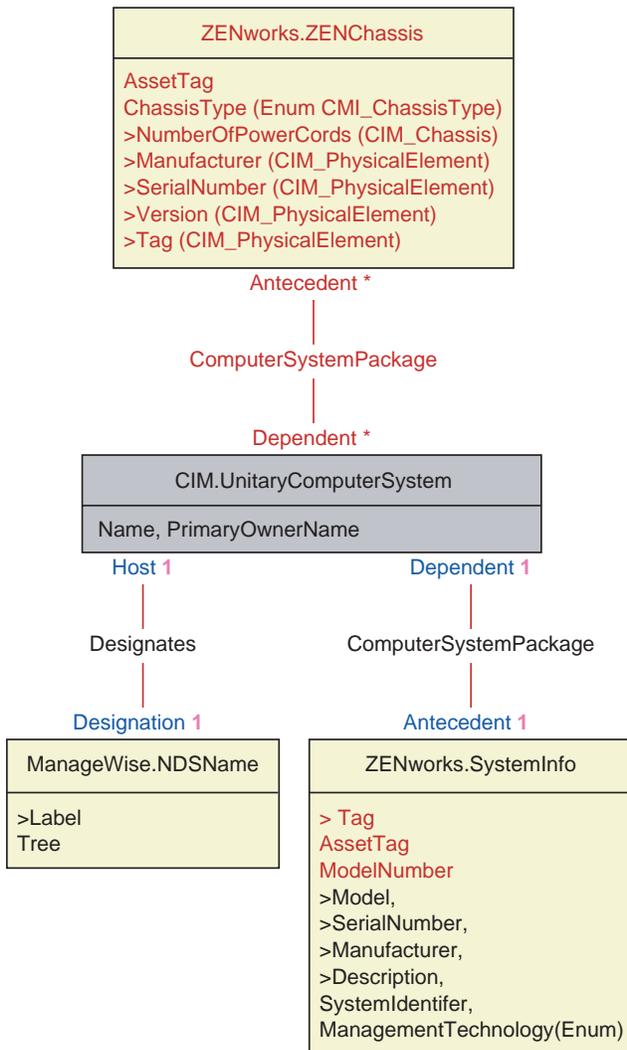


Figure 15-8 Schema for Monitor

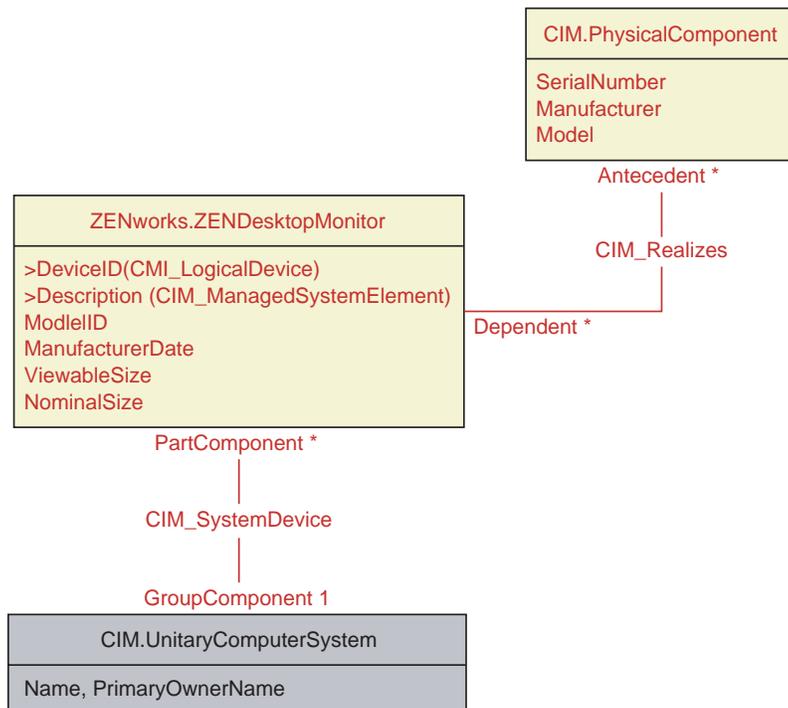


Figure 15-9 Schema for Input devices, Port, Driver, User information, and BIOS

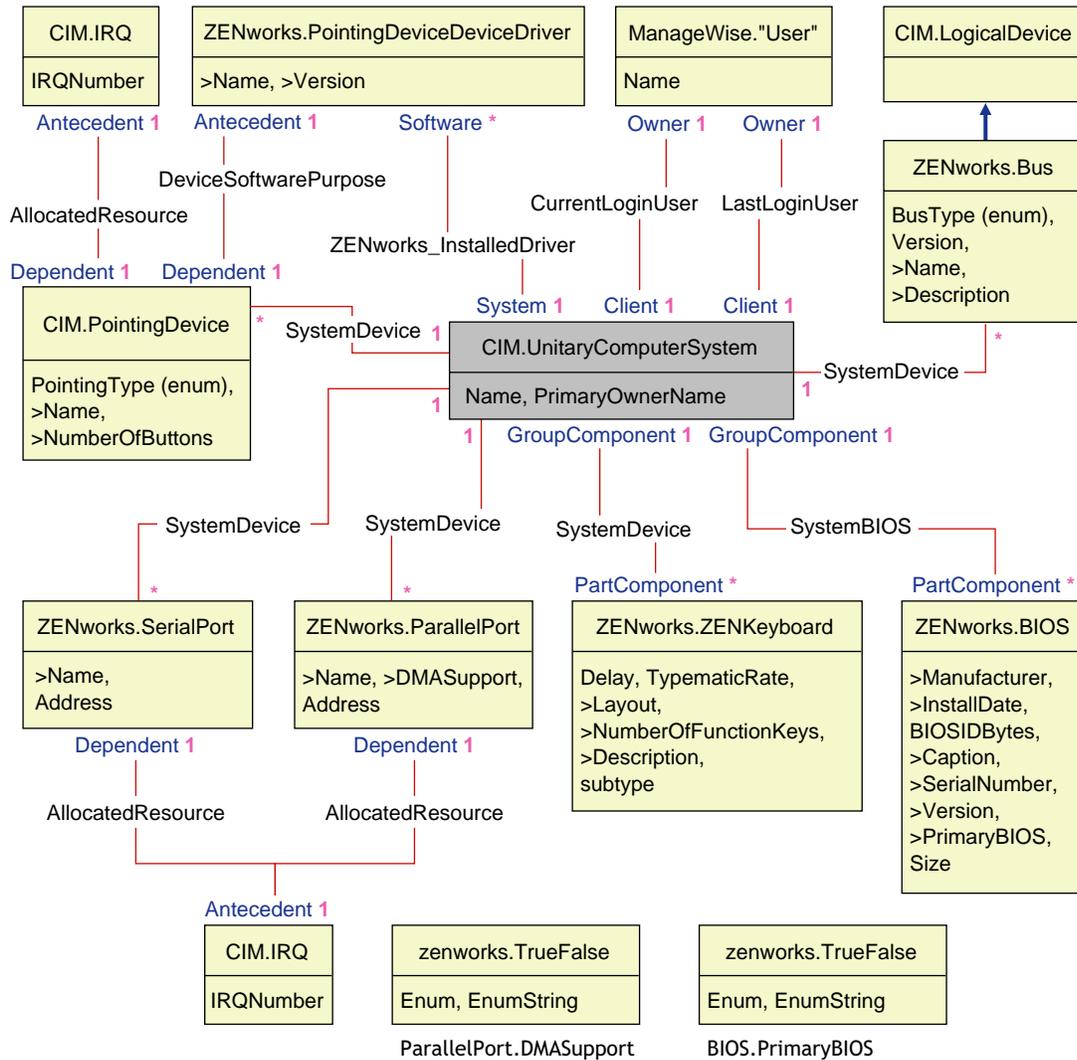


Figure 15-10 Schema for Storage Media

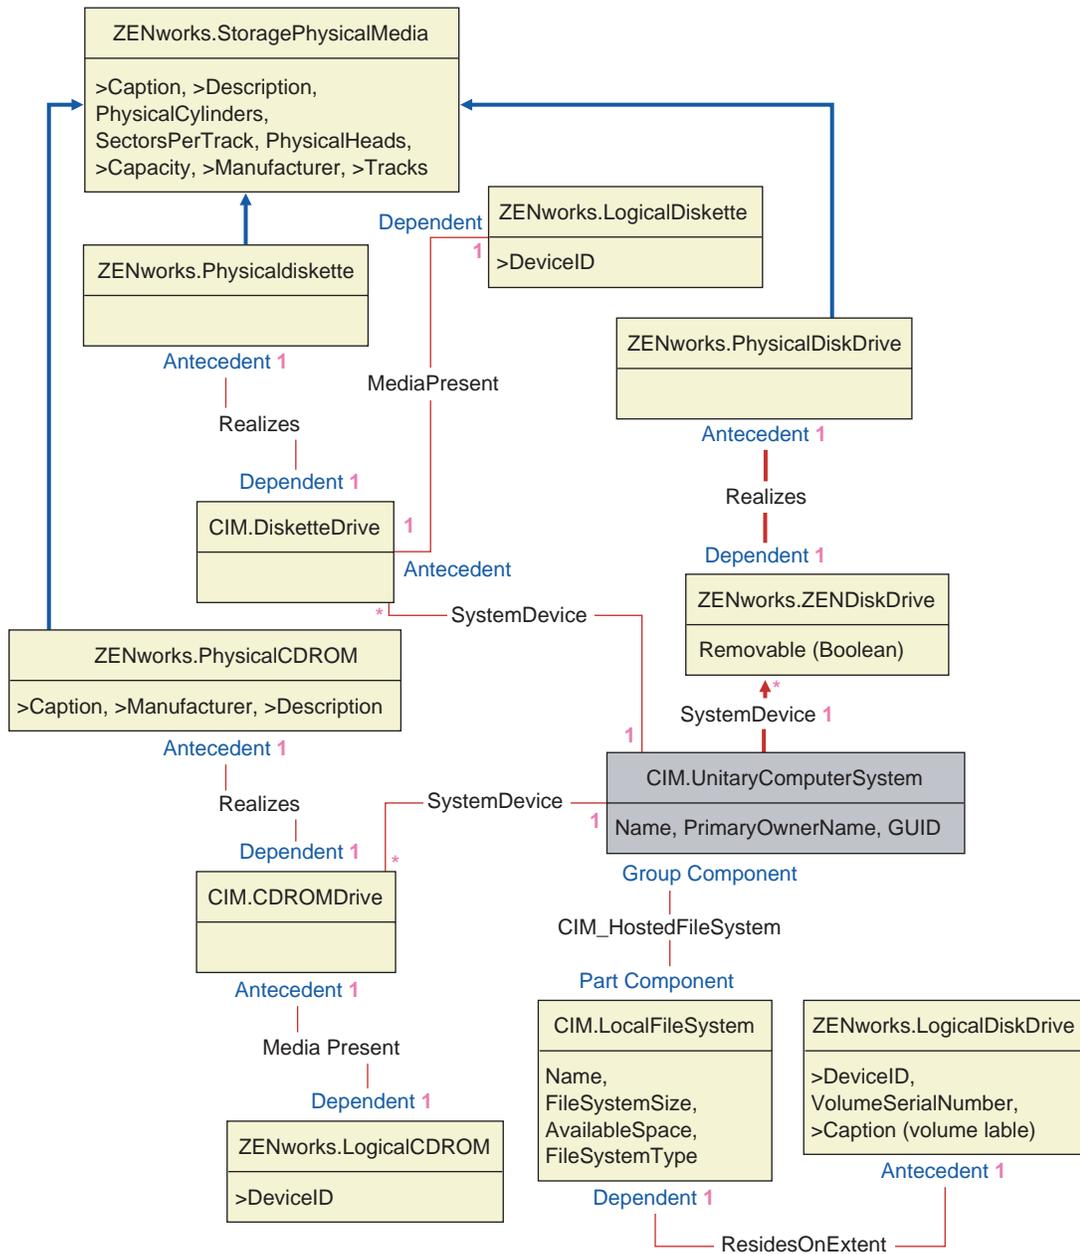


Figure 15-11 Schema for Network, Modem, and Sound Adapter

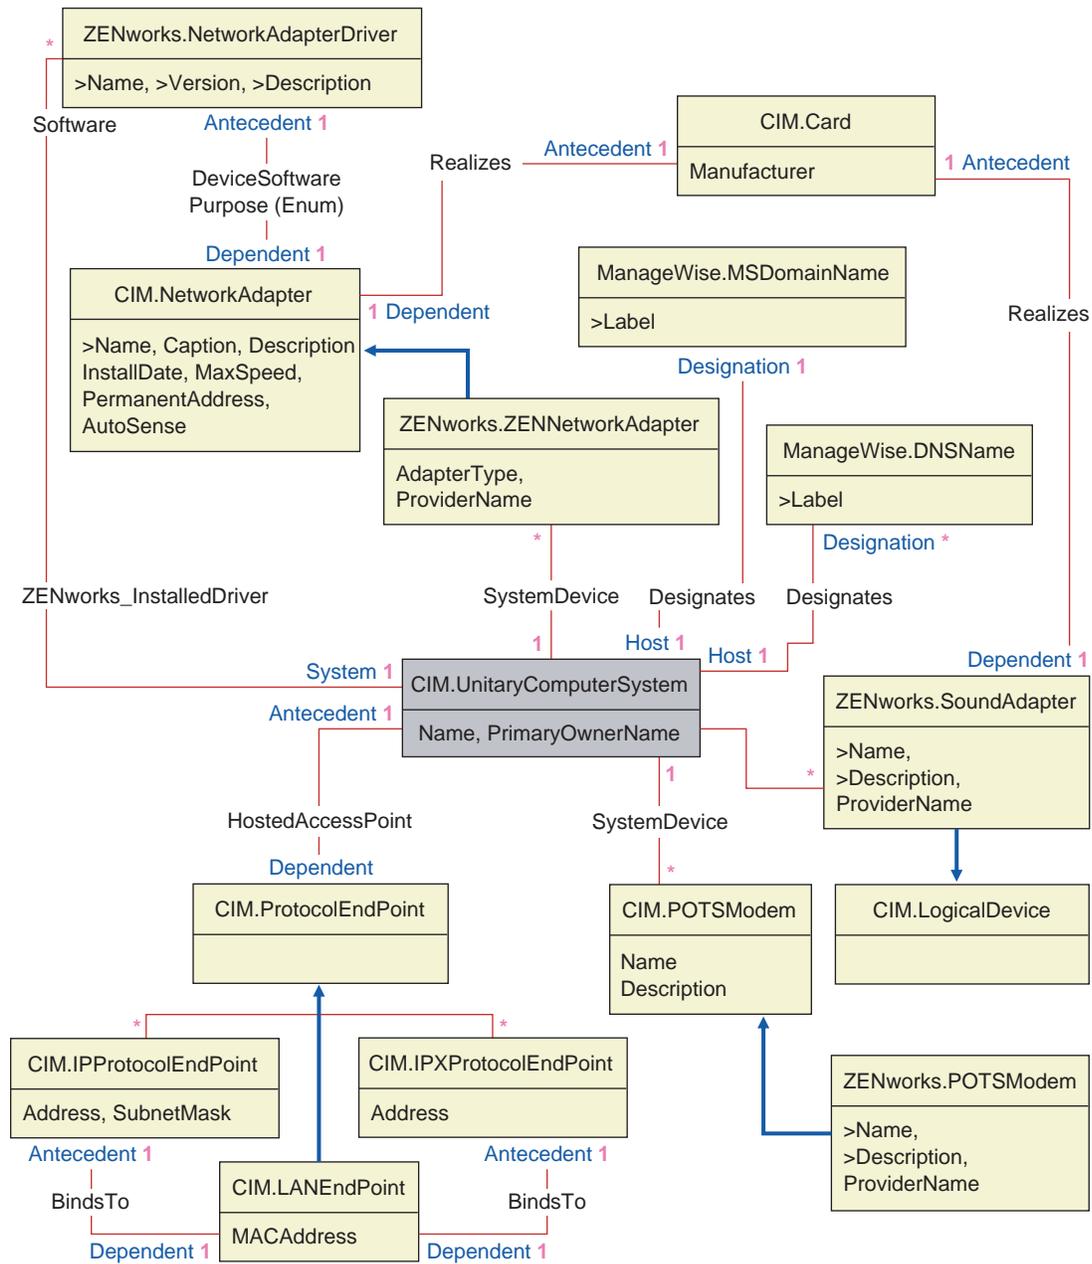


Figure 15-12 Schema for Battery, Card, Cache, Mother Board, and DMA

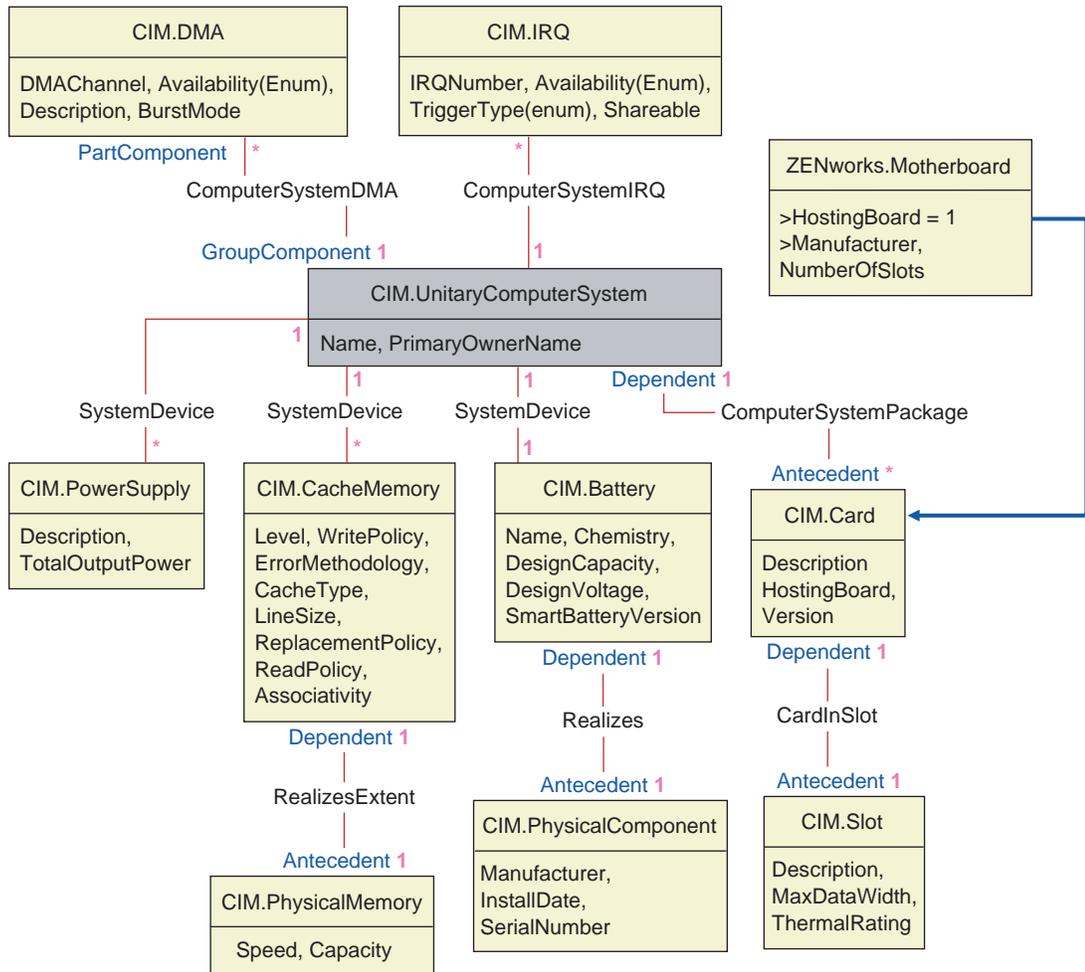
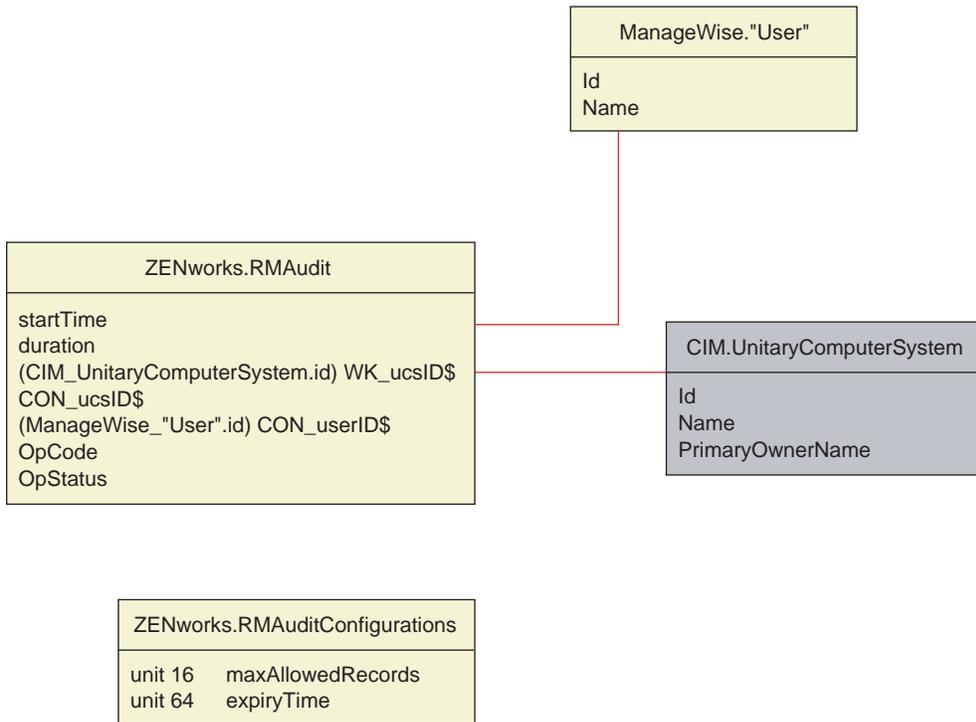


Figure 15-13 Schema for Remote Management Audit

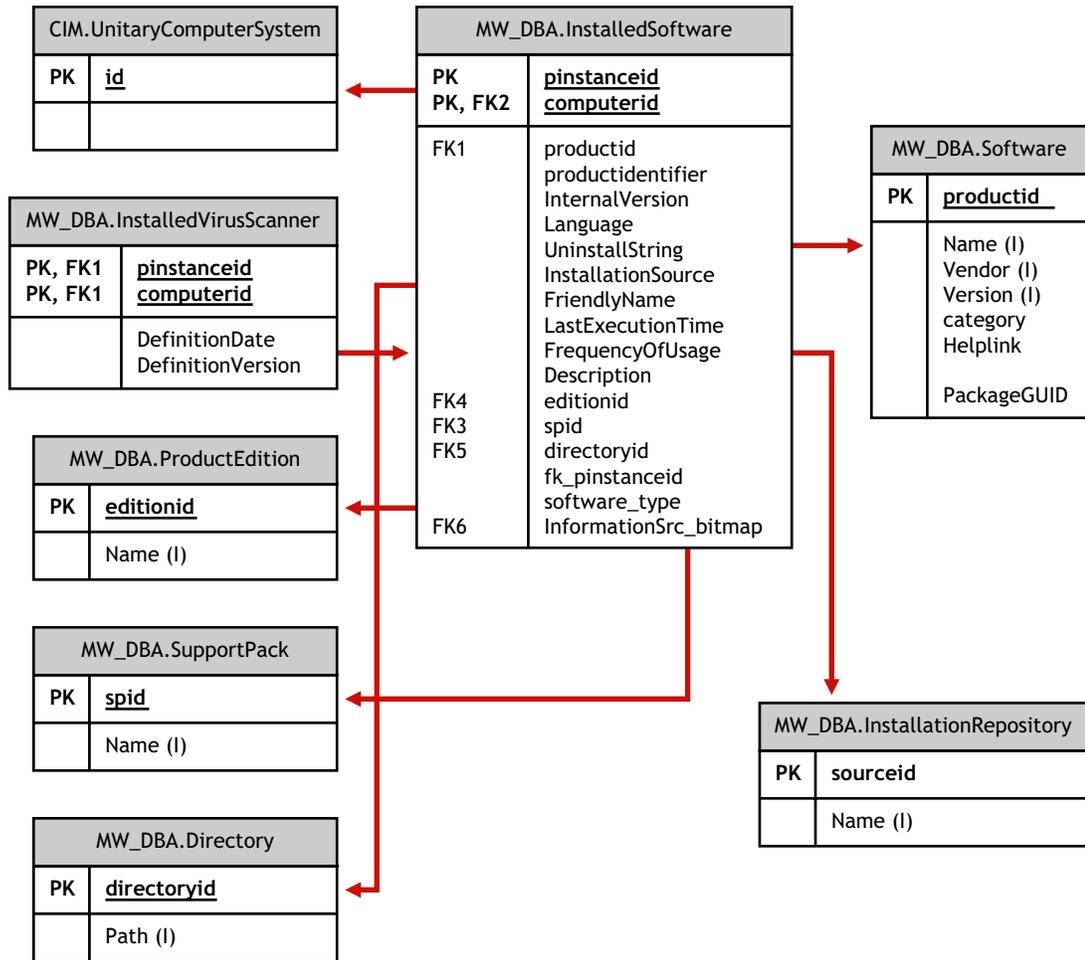


15.3.4 Software Inventory Schema

The following software schema diagrams model the Inventory database in ZENworks 7 Server Management. In the following schema diagram, the `CIM.UnitaryComputerSystem` represents a managed inventory system.

For more information about the tables, see [Appendix L, "ZENworks 7 Server Management Inventory Attributes,"](#) on page 747.

Figure 15-14 Software Inventory Schema Diagram 1

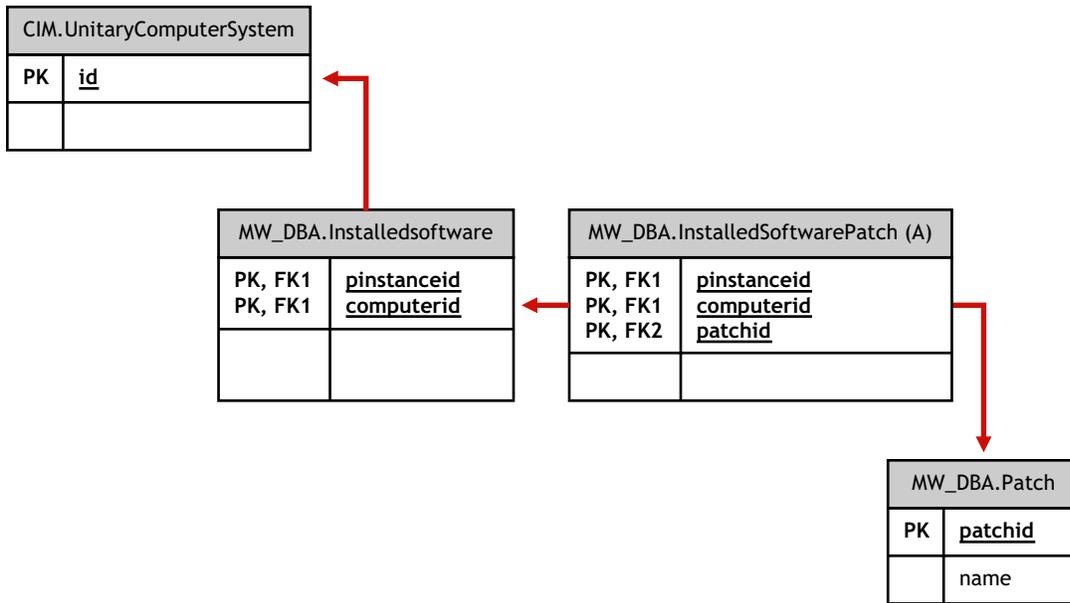


In Figure 15-14, class MW_DBA.Software associates to CIM.UnitaryComputerSystem using the association MW_DBA.InstalledSoftware with MW_DBA.InstalledSoftware.ComputerSystem pointing to CIM.UnitaryComputerSystem and MW_DBA.InstalledSoftware.ProductID pointing to MW_DBA.Software. The relationship between the two classes is one to many. This means a computer system might have more than one software information.

MW_DBA.InstalledSoftware association has Foreign key references to the following tables: ProductEdition, SupportPack, Directory, and Installation Repository.

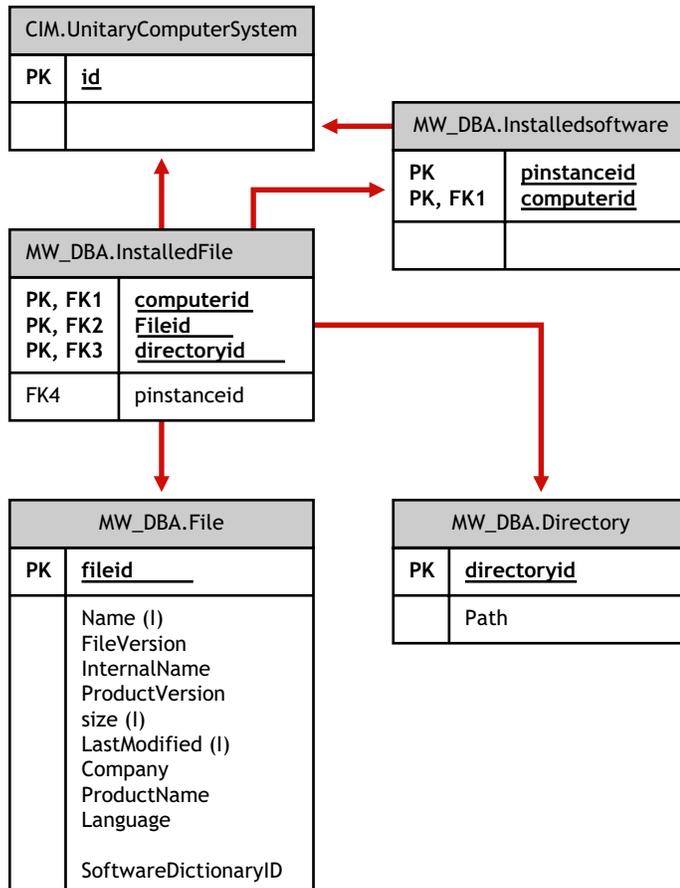
MW_DBA.InstalledVirusScanner inherits the software information from MW_DBA.InstalledSoftware along with virus specific information such as Definition date and Definition version.

Figure 15-15 Software Patch Inventory Schema Diagram 2



In Figure 15-15, class MW_DBA.Patch associates to MW_DBA.InstalledSoftware using the association MW_DBA.InstalledSoftwarePatch with MW_DBA.InstalledSoftwarePatch.instanceID pointing to MW_DBA.InstalledSoftware and MW_DBA.InstalledSoftwarePatch.PatchID pointing to MW_DBA.Patch. The relationship between the two classes is one to many. This means a software might have zero or more patch information.

Figure 15-16 Schema for File and Directory Information



In **Figure 15-16**, class MW_DBA.File associates to MW_DBA.InstalledSoftware using the association MW_DBA.InstalledFile with MW_DBA.InstalledFile.pinstanceID pointing to MW_DBA.InstalledSoftware and MW_DBA.InstalledFile.fileID pointing to MW_DBA.File. The relationship between the two classes is one to many. This means a software might have zero or more file information.

In this illustration, class MW_DBA.Directory associates to MW_DBA.InstalledSoftware using the association MW_DBA.InstalledFile with MW_DBA.InstalledFile.pinstanceID pointing to MW_DBA.InstalledSoftware and MW_DBA.InstalledFile.DirectoryID pointing to MW_DBA.Directory.

Figure 15-17 Schema for Software Sub-classes

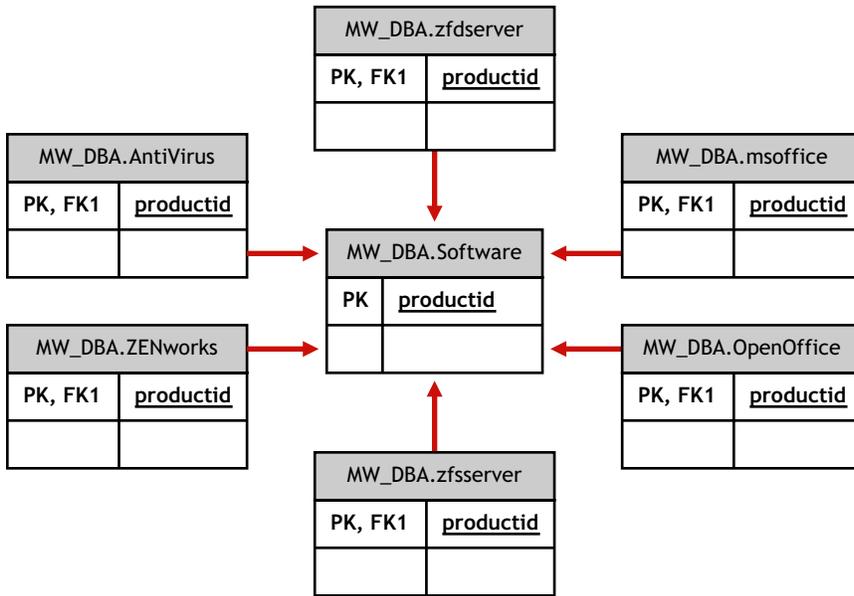


Figure 15-18 Schema for Software Sub-classes

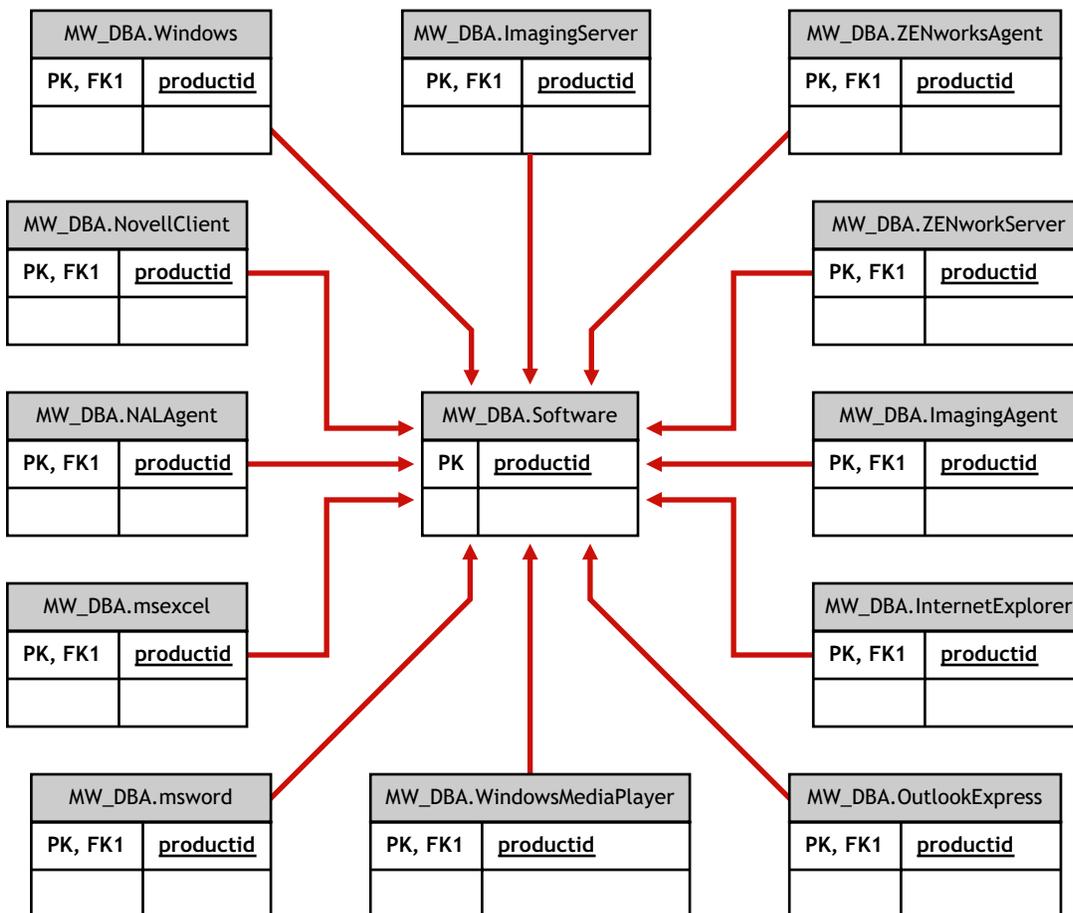
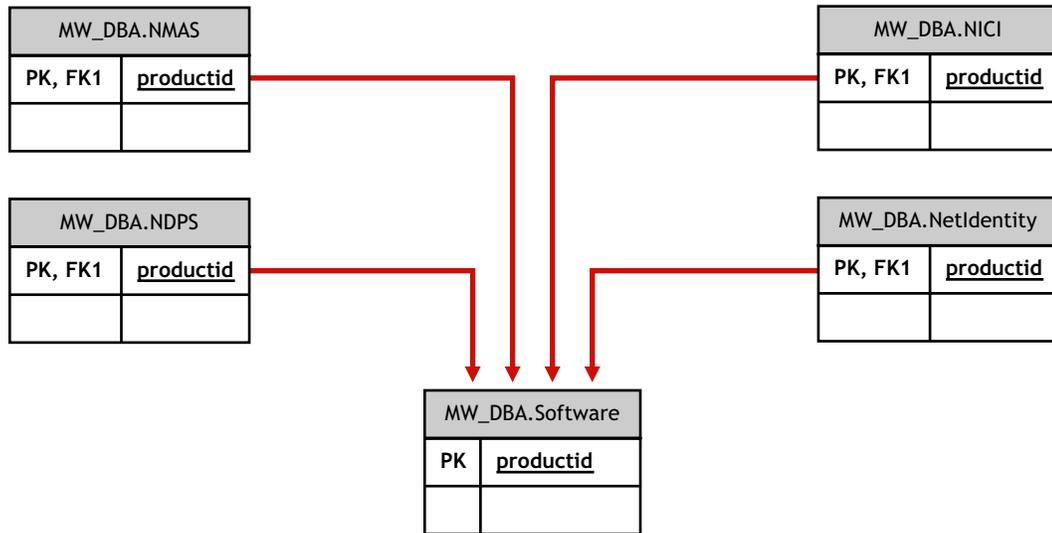


Figure 15-19 Schema for Software Sub-classes



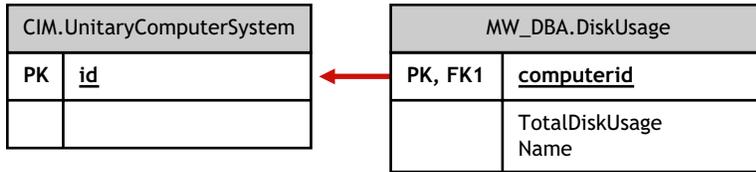
In the above illustrations, MW_DBA.MSoffice inherits the software information from MW_DBA.Software. This sub-class gets directly the MS office information.

This is also applicable for the classes listed in [Table 15-1](#) and [Figure 15-20](#):

Table 15-1 Schema Classes

| | | |
|-------------------------------|---------------------------|----------------------------------|
| mw_dba.zfdserver | mw_dba.zfdinventoryserver | mw_dba.zfdagent |
| mw_dba.zfsserver | mw_dba.zfdinventoryagent | mw_dba.zfsagent |
| mw_dba.zfsinventoryserver | mw_dba.mspowerpoint | mw_dba.msphotodraw |
| mw_dba.zfsinventoryagent | mw_dba.msoutlook | mw_dba.zfdwsmanager |
| mw_dba.zfsrmsserver | mw_dba.msaccess | mw_dba.zfdwsimportserver |
| mw_dba.zfsrmagent | mw_dba.mspublisher | mw_dba.zfdinldbserver |
| mw_dba.zfdrmsserver | mw_dba.msfrontpage | mw_dba.zfsinldbserver |
| mw_dba.zfdrmagent | mw_dba.msinfopath | mw_dba.zfdinvxmlproxyserver |
| mw_dba.zfsinvxmlproxyserver | mw_dba.zfdimagingagent | mw_dba.zfdimagingserver |
| mw_dba.zfdnalagent | mw_dba.zfdnalserver | mw_dba.zfdnaldb |
| mw_dba.middletier | mw_dba.zfsmssserver | mw_dba.zfspds |
| mw_dba.zfspxeserver | mw_dba.zfsmssrvgmtagent | mw_dba.zfsmstrafficanalysisagent |
| mw_dba.zfsmmsadctrendingagent | mw_dba.zfspdsdb | mw_dba.zfhserver |
| mw_dba.zfhaccesspoin | mw_dba.zfhdesktopsync | |

Figure 15-20 Schema for Disk Usage



In the above illustration, MW_DBA.DiskUsage has the computerID column foreign key references to the CIM.UnitaryComputerSystem.ID. The MW_DBA.DiskUsage table contains the total disk usage and the file extension name.

15.3.5 Sample Inventory Database Queries

The following are sample queries for retrieving the inventory information from the ZENworks 7 Server Management Inventory database.

Refer to the schema diagrams in [“Schema Diagrams of CIM and the Extension Schema in ZENworks 7 Server Management”](#) on page 559 to find out the associated schema classes and attribute information.

1. Retrieve the name and ID of all inventoried servers from the database and also to the eDirectory tree to which these servers are registered. The query is as follows:

```
SELECT
    u.id$, u.name, m.tree
FROM
    ManageWise.NDSName m,
    CIM.UnitaryComputerSystem u,
    ManageWise.Designates s
WHERE
    s.Designation=m.id$ AND s.Host=u.id$;
```

In the above query, the tree name is part of the computer system name.

2. Retrieve the asset tag, manufacturer, and model number of all the inventoried servers in the database. The query is as follows:

```
SELECT
    m.AssetTag,
    m.Manufacturer,
    m.ModelNumber,
    m.SerialNumber
FROM
    CIM.UnitaryComputerSystem u,
    CIM.ComputerSystemPackage s,
    ZENworks.SystemInfo m
WHERE
    s.Antecedent=m.id$ AND s.Dependent=u.id$;
```

3. Retrieve all the Microsoft applications with their versions and IDs that are installed on the inventoried server 'SJOHN164_99_139_79' registered under the NOVELL_AUS eDirectory tree. The query is as follows:

```

SELECT
    m.Name,
    m.Version,
    im.ProductIdentifier
FROM
    CIM.UnitaryComputerSystem u,
    MW_DBA.InstalledSoftware im,
    MW_DBA.Software m
WHERE
    u.Name='SJOHN164_99_139_79.Novell_AUS' AND
    (im.computerid=u.id$ and im.productid=m.productid)
    AND m.Vendor LIKE 'Microsoft%';

```

4. Retrieve the processor information for the inventoried server 'SJOHN164_99_139_79.NOVELL_AUS'. The query is as follows:

```

SELECT
    procr.DeviceID,
    role.EnumString,
    family.EnumString,
    procr.OtherFamilyDescription,
    upg.EnumString,
    procr.MaxClockSpeed,
    procr.CurrentClockSpeed,
    procr.Stepping
FROM
    CIM.UnitaryComputerSystem ucs,
    CIM.ComputerSystemProcessor csp,
    CIM.Processor procr,
    CIM.Role_en_US role,
    CIM.Family_en_US family,
    CIM.UpgradeMethod_en_US upg
WHERE
    ucs.name='SJOHN164_99_139_79.Novell_AUS' AND
    csp.PartComponent=procr.id$ AND
    (
        (
            ( procr.Role IS NOT NULL AND procr.Role=role.Enum ) OR
            ( procr.Role IS NULL AND role.Enum=1000 )
        )
        AND
        procr.Family=family.Enum
        AND
        (
            ( procr.UpgradeMethod IS NOT NULL AND
            procr.UpgradeMethod=upg.Enum ) OR
            ( procr.UpgradeMethod IS NULL AND upg.Enum=1000 )
        )
    )

```

```
);
```

5. Retrieve the ID of the UnitaryComputerSystem used for the inventoried server 'SJOHN164_99_139_79.NOVELL_AUS'. The query is as follows:

```
SELECT
    id$
FROM
    CIM.UnitaryComputerSystem
WHERE
    Name='SJOHN164_99_139_79.Novell_AUS';
```

6. Find the number of inventoried servers in the database. The query is as follows:

```
SELECT
    count(u.id$)
FROM
    CIM.UnitaryComputerSystem u,
    CIM.InstalledSoftwareElement s,
    ZENworks.InventoryScanner m
WHERE
    m.id$=s.Software AND u.id$=s.System;
```

7. When you know the ID of the UnitaryComputerSystem for a particular inventoried server from the query as shown in query 5, query 4 can be modified as:

```
SELECT
    procr.DeviceID,
    role.EnumString,
    family.EnumString,
    procr.OtherFamilyDescription,
    upg.EnumString,
    procr.MaxClockSpeed,
    procr.CurrentClockSpeed,
    procr.Stepping
FROM
    CIM.UnitaryComputerSystem ucs,
    CIM.ComputerSystemProcessor csp,
    CIM.Processor procr,
    CIM.Role_en_US role,
    CIM.Family_en_US family,
    CIM.UpgradeMethod_en_US upg
WHERE
    ucs.id$ = ? AND
    csp.PartComponent=procr.id$ AND
    (
        (
            ( procr.Role IS NOT NULL AND procr.Role=role.Enum ) OR
            ( procr.Role IS NULL AND role.Enum=1000 )
        )
    )
    AND
    procr.Family=family.Enum
```

```

        AND
        (
            ( procr.UpgradeMethod IS NOT NULL AND
procr.UpgradeMethod=upg.Enum ) OR
            ( procr.UpgradeMethod IS NULL AND upg.Enum=1000 )
        )
    );

```

Substitute the ID of the specified inventoried server in place of the ?, value for ucs.id\$ in the query.

8. List the IP address, IPX address, and MAC address of all servers in the database. The query is as follows:

```

SELECT
    u.name,
    ip.Address,
    ipx.Address,
    mac.MACAddress
FROM
    CIM.UnitaryComputerSystem u,
    CIM.HostedAccessPoint s1,
    CIM.IPProtocolEndpoint ip,
    CIM.HostedAccessPoint s2,
    CIM.IPXProtocolEndpoint ipx,
    CIM.HostedAccessPoint s3,
    CIM.LANEndpoint mac
WHERE
    (s1.Dependent=ip.id$ and s1.Antecedent=u.id$) AND
    (s2.Dependent=ipx.id$ and s2.Antecedent=u.id$) AND
    (s3.Dependent=mac.id$ and s3.Antecedent=u.id$);

```

9. Retrieve the name and other properties of the drives on the hard disk of the specified inventoried server. The query is as follows:

```

SELECT
    n.Name,
    m.DeviceID,
    n.FileSystemSize,
    n.AvailableSpace,
    n.FileSystemType,
    m.VolumeSerialNumber,
    m.caption as VolumeLabel
FROM
    CIM.HostedFileSystem s,
    CIM.LocalFileSystem n,
    CIM.ResidesOnExtent r,
    ZENworks.LogicalDiskDrive m
WHERE
    (s.GroupComponent=? and s.PartComponent=n.id$) AND
    (r.Dependent=n.id$ and r.Antecedent=m.id$);

```

10. Retrieve all Custom attribute information stored in the database. The query is as follows:

```
SELECT * FROM ZENworks.CustomInformation;
```

11. Retrieve all Custom attribute information associated to the Class CIM.UnitaryComputerSystem. The query is as follows:

```
SELECT
  *
FROM
  ZENworks.CustomInformation
WHERE
  extractClass(id) IN
  (SELECT id FROM MW_DBA.t$Class WHERE
   ClassName='CIM. UnitaryComputerSystem')
```

12. Retrieve all the Microsoft Office installations in the enterprise. The query is as follows:

```
SELECT
  u.name,
  m.FriendlyName,
  im.InternalVersion,
  im.ProductIdentifier
FROM
  CIM.UnitaryComputerSystem u,
  MW_DBA.InstalledSoftware im,
  MW_DBA.Software m,
  MW_DBA.MSOffice mso
WHERE
  mso.id$=m.productid AND
  m.productid=im.productid AND
  im.computerid=u.id$;
```

13. Retrieve all the Internet Explorer installations in the enterprise. The query is as follows:

```
SELECT
  u.Name,
  m.Name,
  m.Version,
  im.InternalVersion,
  im.ProductIdentifier
FROM
  CIM.UnitaryComputerSystem u,
  MW_DBA.InstalledSoftware im,
  MW_DBA.Software m,
  MW_DBA.InternetExplorer ie
WHERE
  ie.id$=m.productid AND
  m.productid=im.productid AND
  im.computerid=u.id$;
```

NOTE: Query 12 and 13 follow nearly the same syntax except for the table relating to the component. A similar approach can be used for the components such as Windows Media

Player, Outlook Express, Microsoft Word, and Microsoft Excel. The complete set of these tables is available in the Schema.

14. Retrieve all the Anti-Virus installations in the enterprise. The query is as follows:

```
SELECT
    u.Name,
    m.Name,
    m.Version,
    im.InternalVersion,
    ivs.DefinitionVersion,
    ivs.DefinitionDate
FROM
    CIM.UnitaryComputerSystem u,
    MW_DBA.InstalledSoftware im,
    MW_DBA.Software m,
    MW_DBA.InstalledVirusScanner ivs
WHERE
    ivs.pinstanceid=im.pinstanceid AND
    m.productid=im.productid AND
    im.computerid=u.id$;
```

15. Retrieve all the applications and the details of the files associated with the application that are installed on the inventoried server 'SJOHN164_99_139_79.NOVELL_AUS'. The query is as follows:

```
SELECT
    u.Name,
    m.Name,
    m.Version,
    m.Category,
    zfile.company,
    zfile.productname,
    zfile.productversion,
    zfile.name,
    dir.path,
    zfile.fileversion,
    zfile."size",
    zfile.lastmodified,
    zfile.internalname,
    zfile.softwaredictionaryid
FROM
    CIM.UnitaryComputerSystem u,
    MW_DBA.InstalledSoftware iso,
    MW_DBA.Software m,
    MW_DBA.InstalledFile ifile,
    MW_DBA."file" zfile,
    MW_DBA.Directory dir
WHERE
    u.Name=' SJOHN164_99_139_79.Novell_AUS' AND
```

```

iso.computerid=u.id$ AND
iso.productid=m.productid AND
iso.pinstanceid=ifile.pinstanceid AND
    ifile.directoryid=dir.id AND
    ifile.fileid=zfile.id;

```

16. Retrieve all the files present on the inventoried server 'SJOHN164_99_139_79.NOVELL_AUS' which has not been associated with a valid software. The query is as follows:

```

SELECT
    u.Name,
    zfile.name,
    dir.path,
    zfile.fileversion,
    zfile."size",
    zfile.lastmodified,
    zfile.internalname,
    zfile.productversion,
    zfile.company,
    zfile.productname
FROM
    CIM.UnitaryComputerSystem u,
    MW_DBA.InstalledFile ifile,
    MW_DBA."file" zfile,
    MW_DBA.Directory dir
WHERE
    u.Name=' SJOHN164_99_139_79.Novell_AUS' AND
    u.id$=ifile.computerid AND
    ifile.fileid=zfile.id AND
    ifile.directoryid=dir.id AND
    ifile.pinstanceid is null;

```

17. Retrieve the disk usage details of files with known extensions on each inventoried machine in the enterprise. The query is as follows:

```

SELECT
    u.Name,
    du.Name,
    du.TotalDiskUsage
FROM
    CIM.UnitaryComputerSystem u,
    MW_DBA.DiskUsage du
WHERE
    u.id$=du. Computerid AND
    du.Name is not null;

```


Managing Your Inventory Information

16

This section contains the following information to help you customize the way Novell® ZENworks® 7 Server Inventory displays information:

- ◆ Section 16.1, “Viewing the Inventory Servers Deployed for Inventory,” on page 581
- ◆ Section 16.2, “Customizing the Hardware Inventory Information To Be Scanned,” on page 582
- ◆ Section 16.3, “Customizing the Software Inventory Information To Be Scanned For the ZENworks 7 Inventoried Servers,” on page 586
- ◆ Section 16.4, “Customizing the Software Inventory Information To Be Scanned For ZENworks for Servers 3.x Inventoried Servers,” on page 632
- ◆ Section 16.5, “Removing Redundant Inventoried Servers from the Inventory Database,” on page 632

16.1 Viewing the Inventory Servers Deployed for Inventory

Using ConsoleOne®, you can view the Inventory servers and databases that you configured for collecting inventory.

To get a complete Inventory tree view:

- 1 Log into all the Novell eDirectory™ trees that contain Inventory servers present in your inventory tree.
- 2 In ConsoleOne, select a container, click the *View* menu, then click *Complete Tree View*.
All the Inventory servers within the container are displayed in the Complete Tree View.

To view a complete tree view if your inventory deployment involves roll-up of information between Inventory servers that are situated on different Novell eDirectory trees:

- 1 In ConsoleOne, select *NDS Tree*.
- 2 Click *View*, then click *Complete Tree View*.
- 3 Select the eDirectory trees or containers within the tree that contains the Inventory servers.
- 4 Click *OK*.

To view all Inventory server from the selected Inventory server to the highest-level server:

- 1 In ConsoleOne, right-click the Inventory Service object (*Inventory Service_server_name*), click *View*, then click *Up Tree View* or double-click the Inventory Service object.

If your inventory deployment consists of a single eDirectory tree, an Up Tree View displays all the Inventory servers from the selected Inventory server up to the highest level (Root Server).

If your inventory deployment involves roll-up of inventory information across Inventory servers located on different eDirectory trees, the Up Tree View displays all the Inventory servers from the selected Inventory server up to the highest level server to which you have logged in.

NOTE: You cannot collapse the inventory tree using the short-cut keys.

16.2 Customizing the Hardware Inventory Information To Be Scanned

ZENworks 7 Server Management allows you to collect information that is not part of the default hardware inventory from the inventoried servers.

- ♦ [Section 16.2.1, “Scanning for Vendor-Specific Asset Information from DMI,” on page 582](#)
- ♦ [Section 16.2.2, “Customizing the Hardware Scanning Information of Jaz and Zip Drive Vendors,” on page 583](#)
- ♦ [Section 16.2.3, “Customizing the Hardware Information for Monitor Size,” on page 584](#)

16.2.1 Scanning for Vendor-Specific Asset Information from DMI

- 1 In the Server Inventory policy, click the *Configuration Editor* tab.

For more information, see [Section 13.5, “Configuring the Server Inventory Policy,” on page 522](#).

- 2 Click the *Asset Information* suboption, then click *Set Defaults*.

The following entries are populated.

```
[ASSETTAG]
DMI1_CLASSNAME=
DMI1_ATTRIBUTEID=
DMI2_CLASSNAME=
DMI2_ATTRIBUTEID=
[SERIALNUMBER]
DMI1_CLASSNAME=
DMI1_ATTRIBUTEID=
DMI2_CLASSNAME=
DMI2_ATTRIBUTEID=
[MODEL]
DMI1_CLASSNAME=
DMI1_ATTRIBUTEID=
DMI2_CLASSNAME=
DMI2_ATTRIBUTEID=
[COMPUTERTYPE] DMI1_CLASSNAME=DMI1_ATTRIBUTEID=
[MODELNUMBER] DMI1_CLASSNAME=DMI1_ATTRIBUTEID=
```

- 3 Specify the values.

The Asset Information contains the following sections:

- ♦ Contains Asset Tag in the section [ASSETTAG]

- ◆ Contains Serial Number in the section [SERIALNUMBER]
- ◆ Contains Computer Model in the section [MODEL]
- ◆ Contains Computer Type [COMPUTERTYPE]
- ◆ Contains Computer Model Number [MODELNUMBER]

Each section contains the particular DMI Class name and DMI Class Attribute ID.

The format of Asset Information is as follows:

```
[ASSETTAG]
DMI1_CLASSNAME=DMI_class_name_for_asset_tag
DMI1_ATTRIBUTEID=DMI_attribute_ID_for_asset_tag
[SERIALNUMBER]
DMI1_CLASSNAME=DMI_class_name_for_serial_number
DMI1_ATTRIBUTEID=DMI_attribute_ID_for_serial_number
[MODEL]
DMI1_CLASSNAME=DMI_class_name_for_computer_model
DMI1_ATTRIBUTEID=DMI_attribute_ID_for_computer_model
```

The value of the Asset Information sections can have a maximum string length of 64 characters.

A DMI Class name can be any DMI class other than DMTF|COMPONENTID|00x.

If there is more than one DMI vendor implementing different custom DMI classes, you can specify multiple DMI classes. A maximum of five classes can be specified in these sections. For example, the asset information for five classes is as follows:

```
[ASSETTAG]
DMI1_CLASSNAME=DMI_class_name_for_asset_tag
DMI1_ATTRIBUTEID=DMI_attribute_ID_for_asset_tag
DMI2_CLASSNAME=DMI_class_name_for_asset_tag
DMI2_ATTRIBUTEID=DMI_attribute_ID_for_asset_tag
DMI3_CLASSNAME=DMI_class_name_for_asset_tag
DMI3_ATTRIBUTEID=DMI_attribute_ID_for_asset_tag
DMI4_CLASSNAME=DMI_class_name_for_asset_tag
DMI4_ATTRIBUTEID=DMI_attribute_ID_for_asset_tag
DMI5_CLASSNAME=DMI_class_name_for_asset_tag
DMI5_ATTRIBUTEID=DMI_attribute_ID_for_asset_tag
```

The scanner will process DMI1 and if the values of DMI1 are valid, the scanner will not process the remaining DMI classes.

4 Click *OK*.

5 Run the scans on the inventoried servers.

Verify that the inventory information is in the Inventory Summary window.

16.2.2 Customizing the Hardware Scanning Information of Jaz and Zip Drive Vendors

The scan information of the vendors for devices such as backup and floppy devices is usually unavailable on the inventoried server. Also, if the information is available, the vendor information does not usually contain the details. You can customize and update information about the vendors of

these devices in *Server Inventory policy > Configuration Editor > Zipped Names*. The scanners read this information during the hardware scanning process for these devices.

To customize and update the vendor information for display:

- 1 In the Server Inventory policy, click the *Configuration Editor* tab.

For more information, see [Section 13.5, “Configuring the Server Inventory Policy,” on page 522](#).

- 2 Click the *ZIPPED NAMES* suboption, then click *Set Defaults*.

The default values are displayed.

```
[Identifier]
iomega ZIP 100=Iomega 100MB Backup Device
iomega jaz 1GB=Iomega 1GB Backup Device
IOMEGA ZIP 100 D.13=Iomega Corporation
IOMEGA ZIP 1GB D.13=Iomega Corporation
...
```

The format of each entry in the section is as follows:

```
[Identifier]
device_id=vendor_display_name_you_specify
```

where *device_id* is the unique ID generated and updated in the registry by the vendor during the installation of the device on the inventoried server.

For example, the contents of the section are as follows:

```
[Identifier]
iomega ZIP 100=Iomega 100MB Backup Device
```

This entry is for a 100 MB Zip* drive installed on the inventoried server.

- 3 Add or modify the entries.

If you specify incorrect values for the device ID entry, the device will not be displayed in the Inventory windows.

- 4 Click *OK*.

16.2.3 Customizing the Hardware Information for Monitor Size

The inventory information scanned for a monitor includes the following:

Nominal Size: A number representing the diagonal width of the monitor (the distance from one corner of the screen to the opposite corner of the screen). For example, 17".

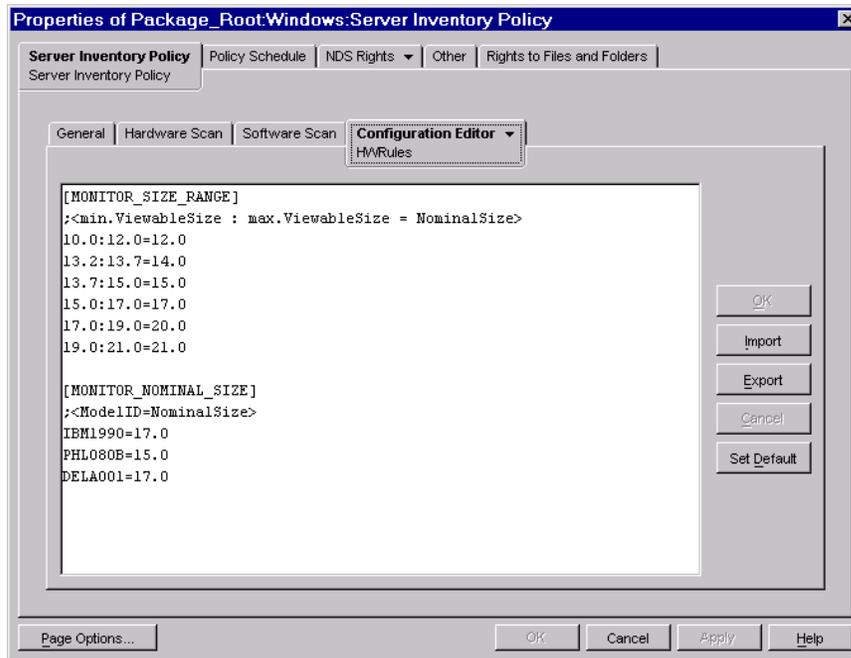
Viewable Size: A number representing the diagonal width of the screen image excluding the black borders around the image's edge. For example, 15.8".

The Inventory scanner automatically scans for the viewable size of the monitor of Windows* inventoried servers. Using the Server Inventory policy, you can customize the nominal size of the monitor to be reported.

IMPORTANT: The Inventory scanner reports inventory information for the monitors that are manufactured only after 1997.

To customize the scan of the nominal size of the monitor:

- 1 In the Server Inventory policy, click the *Configuration Editor* tab, then click the *HWRules* suboption.



- 2 Click *Set Default*.

The default values are displayed in the Configuration Editor box.

- 3 Add or modify the entries.

The format of `HWRules.ini` is as follows:

```
[MONITOR_SIZE_RANGE]
minimum_viewable_size_you_specify:
maximum_viewable_size_you_specify = nominal_size_you_specify
[MONITOR_NOMINAL_SIZE]
model_ID_reported_by_scanner = nominal_size_you_specify
```

In the `[MONITOR_SIZE_RANGE]` section, specify the minimum and maximum range of the viewable size, and the corresponding nominal size of the monitor. The Inventory scanner scans for the model ID of the monitor and reports the nominal size configured in the `[MONITOR_NOMINAL_SIZE]` section of the `HWRules.ini` file.

In the `[MONITOR_NOMINAL_SIZE]` section, specify the model ID and its corresponding nominal size as reported by the Inventory scanner. This information is available in the Inventory Summary dialog box under the Hardware/Software Inventory > Hardware > Monitor attribute.

If the scanned model ID is not listed in `[MONITOR_NOMINAL_SIZE]`, then the scanner scans the viewable size of the monitors. On the basis of the viewable size, the scanner reports the nominal size configured in the `[MONITOR_SIZE_RANGE]` section of the `HWRules.ini` file.

For example, the contents of the `HWRules.ini` file could be as follows:

```
[MONITOR_SIZE_RANGE]
```

```
10.0:12.0=12.0
13.2:13.7=14.0
[MONITOR_NOMINAL_SIZE]
IBM1990=17.0
PHL080B=15.0
```

4 Click *OK* to save the contents to the Server Inventory policy.

16.3 Customizing the Software Inventory Information To Be Scanned For the ZENworks 7 Inventoried Servers

The Inventory scanner has been enhanced with the following features that enable you to control the scanning process more effectively and efficiently for inventoried servers having ZENworks 7 Server Inventory:

- ◆ It supports scanning for the following software inventory information:
 - ◆ Windows operating system and its patches
 - ◆ Internet Explorer and its patches
 - ◆ Windows Media Player and its patches
 - ◆ Outlook Express and its patches
 - ◆ Novell Client32™ and its installed components
 - ◆ ZENworks suite and its installed components
 - ◆ Microsoft Office and its installed applications
 - ◆ Antivirus products such as Symantec Antivirus Corporate Edition and McAfee Antivirus
 - ◆ Virus definition date and version for the antivirus products such as Symantec Antivirus Corporate Edition and McAfee Antivirus
- ◆ It supports scanning for the products listed in the Windows Add/Remove Programs and the MSI database.
- ◆ It includes dictionary of software titles to provide more accurate report of Installed software.
- ◆ It provides rules to control the scope of software scan.
- ◆ It reports total disk usage against configured file extensions.

This section provides information on the following topics:

- ◆ [Section 16.3.1, “What is the ZENworks Software Dictionary?,” on page 587](#)
- ◆ [Section 16.3.2, “What is a Software Dictionary Rule?,” on page 588](#)
- ◆ [Section 16.3.3, “What is a Software Identifier?,” on page 588](#)
- ◆ [Section 16.3.4, “What is a Key Identifier?,” on page 588](#)
- ◆ [Section 16.3.5, “What is an Unidentified Software?,” on page 588](#)
- ◆ [Section 16.3.6, “What is an Inherited Rule?,” on page 588](#)
- ◆ [Section 16.3.7, “What is An Overriding Rule?,” on page 588](#)
- ◆ [Section 16.3.8, “Understanding the Usage and Precedence of ZENworks Software Dictionary Rules,” on page 589](#)

- ◆ Section 16.3.9, “Understanding the Software Dictionary Pattern Types,” on page 595
- ◆ Section 16.3.10, “Configuring the Software Dictionary Rules,” on page 596
- ◆ Section 16.3.11, “Ignore Default File-Software Mapping Rules,” on page 599
- ◆ Section 16.3.12, “Software Dictionary,” on page 599
- ◆ Section 16.3.13, “Report Files with These File Extensions As Unidentified Software,” on page 603
- ◆ Section 16.3.14, “Manage Unidentified Software,” on page 604
- ◆ Section 16.3.15, “Report Multiple Software Versions,” on page 605
- ◆ Section 16.3.16, “Report Disk Space Used by File Extensions,” on page 607
- ◆ Section 16.3.17, “Software Scanning Filters - Drives and Directories,” on page 608
- ◆ Section 16.3.18, “Software Scanning Filters - File Extensions,” on page 613
- ◆ Section 16.3.19, “Software Scanning Filters - Files,” on page 615
- ◆ Section 16.3.20, “Software Scanning Filters - Software,” on page 616
- ◆ Section 16.3.21, “Disk Usage Scanning Filters - Drives and Directories,” on page 618
- ◆ Section 16.3.22, “Disk Usage Scanning Filters - Files,” on page 622
- ◆ Section 16.3.23, “Vendor Name Aliases,” on page 623
- ◆ Section 16.3.24, “Software Name Aliases,” on page 625
- ◆ Section 16.3.25, “Reconcile Software,” on page 626
- ◆ Section 16.3.26, “Sorting Entries in the Table,” on page 627
- ◆ Section 16.3.27, “Filtering Entries in the Table,” on page 627
- ◆ Section 16.3.28, “Refreshing Entries in the Table,” on page 628
- ◆ Section 16.3.29, “Disabling File Scan,” on page 628
- ◆ Section 16.3.30, “Base-lining the Software Dictionary Deployment,” on page 629
- ◆ Section 16.3.31, “Viewing Software Information in the Inventory Summary,” on page 630
- ◆ Section 16.3.32, “Generating Software Inventory Reports,” on page 630

16.3.1 What is the ZENworks Software Dictionary?

The ZENworks software dictionary contains a list of software identifiers and rules. Each software identifier identifies a particular product installed on an inventoried server. The rules control the scope of the scanning process.

The ZENworks software dictionary is automatically installed on an Inventory Server and inventoried servers when you install the Server Inventory software. After you configure the required policies and start the Inventory service, the Inventory scanner reports the software information on the basis of the software dictionary.

There are two types of ZENworks software dictionary: General dictionary and Private dictionary.

General Dictionary: The General dictionary is the part of the software dictionary that contains predefined software identifiers. On the basis of this dictionary, the Inventory scanner reports whether a particular product is installed on an inventoried server.

Private Dictionary: The private dictionary is the part of the software dictionary that contains user-defined software identifiers and rules that enable you to define the scope of Inventory scan and customize the software information. You can configure the rules. For more information on how to configure the rules, see [Section 16.3.10, “Configuring the Software Dictionary Rules,” on page 596.](#)

IMPORTANT: The rules that you define in the private dictionary overrides the predefined rules in the general dictionary.

16.3.2 What is a Software Dictionary Rule?

A software dictionary rule represents a set of conditions that control the scope of scanning process.

16.3.3 What is a Software Identifier?

An entry that identifies a software product is called as software identifier. Each software identifier has a set of file matching attributes and corresponding software information attributes. During the Inventory scan, the scanner reads the attributes from the file headers, and if these attributes match the attributes configured in the dictionary, the information in the corresponding software information attributes is stored in the Inventory database.

16.3.4 What is a Key Identifier?

A software product might be identified through more than one software identifier in the dictionary. In such a scenario, the inventory scanner arbitrarily selects the software information from one of these software identifiers. A key identifier identifies the software identifier from which the inventory scanner should select the software information. The key identifier is useful when the different software identifiers have marginal differences between the values of the attributes (such as Description) and you want the inventory scanner to select the information from a specific software identifier.

16.3.5 What is an Unidentified Software?

An unidentified software has the following characteristics:

- ♦ It is installed on the inventoried servers.
- ♦ It is configured in the [Report Files with These File Extensions As Unidentified Software](#) rule in ZENworks software dictionary.
- ♦ It is not configured in the [Software Dictionary](#) table.

16.3.6 What is an Inherited Rule?

An inherited rule is an entry in the software dictionary that is obtained from another Inventory server through the dictionary distribution. You cannot edit or delete these rules. These rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

16.3.7 What is An Overriding Rule?

The default software identifier in the General dictionary cannot be modified. But if you want to modify a default software identifier, you must create a new software identifier that overrides the

default identifier. The inventory scanner ignores the default the default identifier in favor of the overridden entry.

To create a software identifier that overrides a default identifier, you must specify same values for all the matching attributes defined in the default identifier and provide new values for the software information attributes.

16.3.8 Understanding the Usage and Precedence of ZENworks Software Dictionary Rules

The ZENworks software dictionary rules follow a precedence order. Some guidelines are applicable to all the software dictionary rules and some guidelines are applicable to certain categories of software dictionary rules. For more information, review the following sections:

- ◆ [“Guidelines Applicable to All Software Dictionary Rules” on page 589](#)
- ◆ [“Precedence between Report Only Maximum Software Version and Report All Software Versions” on page 589](#)
- ◆ [“Precedence of Software Dictionary Rules Grouped in the Software Scanning Category” on page 589](#)
- ◆ [“Precedence of Software Dictionary Rules in the Disk Usage Scanning Category” on page 592](#)

Guidelines Applicable to All Software Dictionary Rules

The following guidelines are applicable to all the software dictionary rules that you configure:

- ◆ All software dictionary rules are applied at the inventoried servers by the inventory scanner.
- ◆ You can change the settings of the software dictionary rules using the Software dictionary ConsoleOne snap-ins. For more information on how to configure the software dictionary rules, see [Section 16.3.10, “Configuring the Software Dictionary Rules,” on page 596](#).
- ◆ Every inventory scan contains the version of dictionary files used for that scan. This information is stored in the inventory database.
- ◆ The user-defined software identifier overrides the default software identifier present in the software dictionary but only one user-defined software identifier can be used at a time to override a default software identifier.

Precedence between Report Only Maximum Software Version and Report All Software Versions

By default, the scanner reports only the highest version of the software installed. If a rule in “Report All Software Versions” conflicts with a rule in “Report Only Maximum Software Version,” then the rule in “Report Only Maximum Software Version” overrides the rule of Report All Software Versions.

Precedence of Software Dictionary Rules Grouped in the Software Scanning Category

The software dictionary rules in the Software Scanning category control the scope of scanning for the files on the local file systems.

The Software Scanning category includes the following software dictionary rules:

- ◆ “Scan File Extensions” on page 614
- ◆ “Ignore File Extensions” on page 614
- ◆ “Scan Directories” on page 612
- ◆ “Ignore Directories” on page 611
- ◆ “Scan Drives” on page 610
- ◆ “Ignore Drives” on page 610
- ◆ Section 16.3.20, “Software Scanning Filters - Software,” on page 616
- ◆ Section 16.3.19, “Software Scanning Filters - Files,” on page 615

If you do not configure any of the rules mentioned above, the Inventory scanner scans for all files on the hard disk of the inventoried servers. If the files have matching software identifiers in the software dictionary, the files are reported as identified software. Otherwise, they are reported as unidentified software.

If you configure the rules mentioned above, they take precedence in the following descending order:

- ◆ Software Scanning Filters - Files
- ◆ Software Scanning Filters - Software
- ◆ Scan File Extensions
- ◆ Ignore File Extensions
- ◆ Scan Directories
- ◆ Ignore Directories
- ◆ Scan Drives
- ◆ Ignore Drives

The following flowcharts illustrate the precedence of these rules:

Figure 16-1 Precedence of Software Dictionary rules in the Software Scanning category

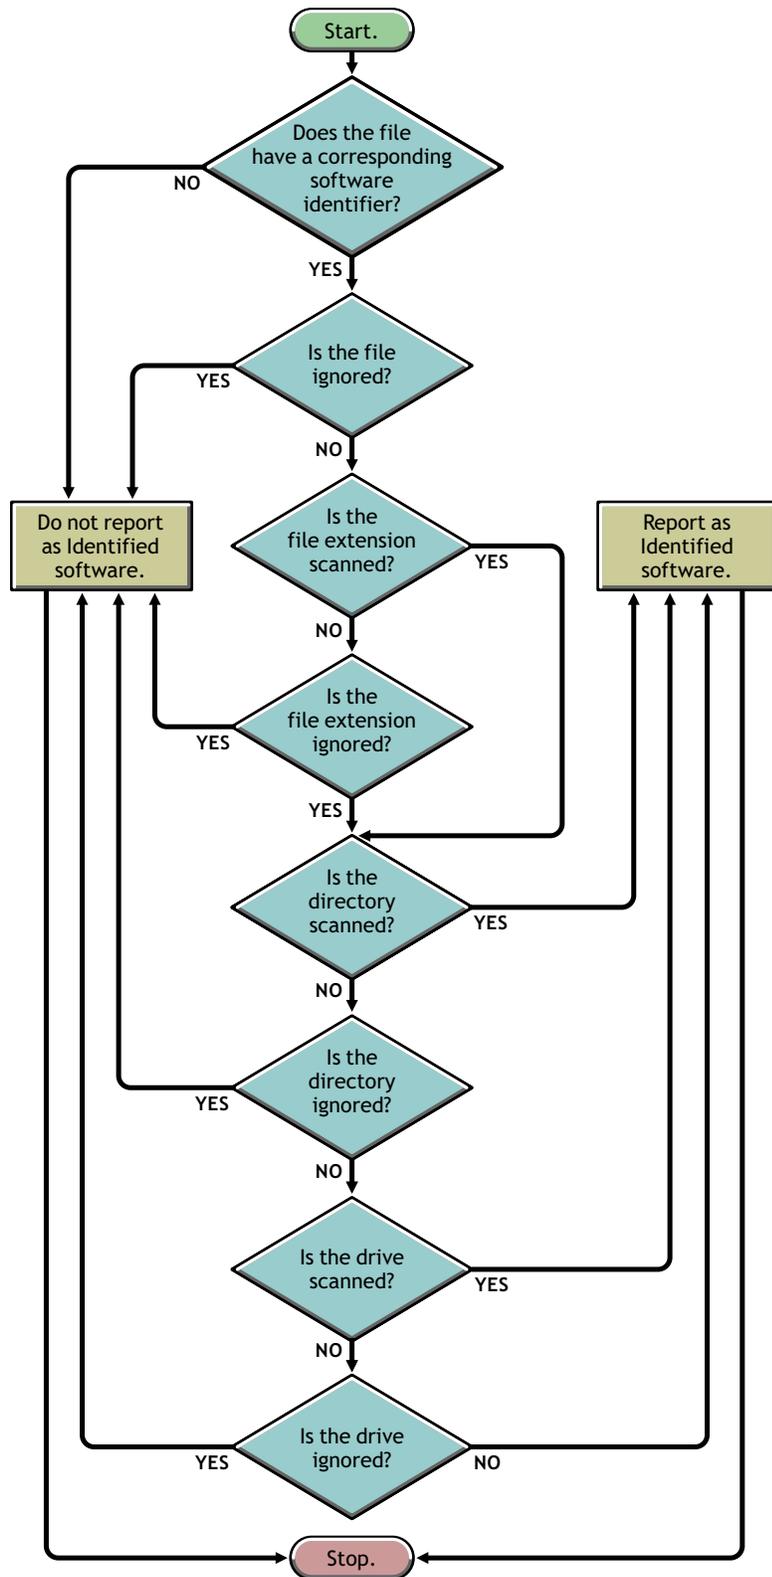
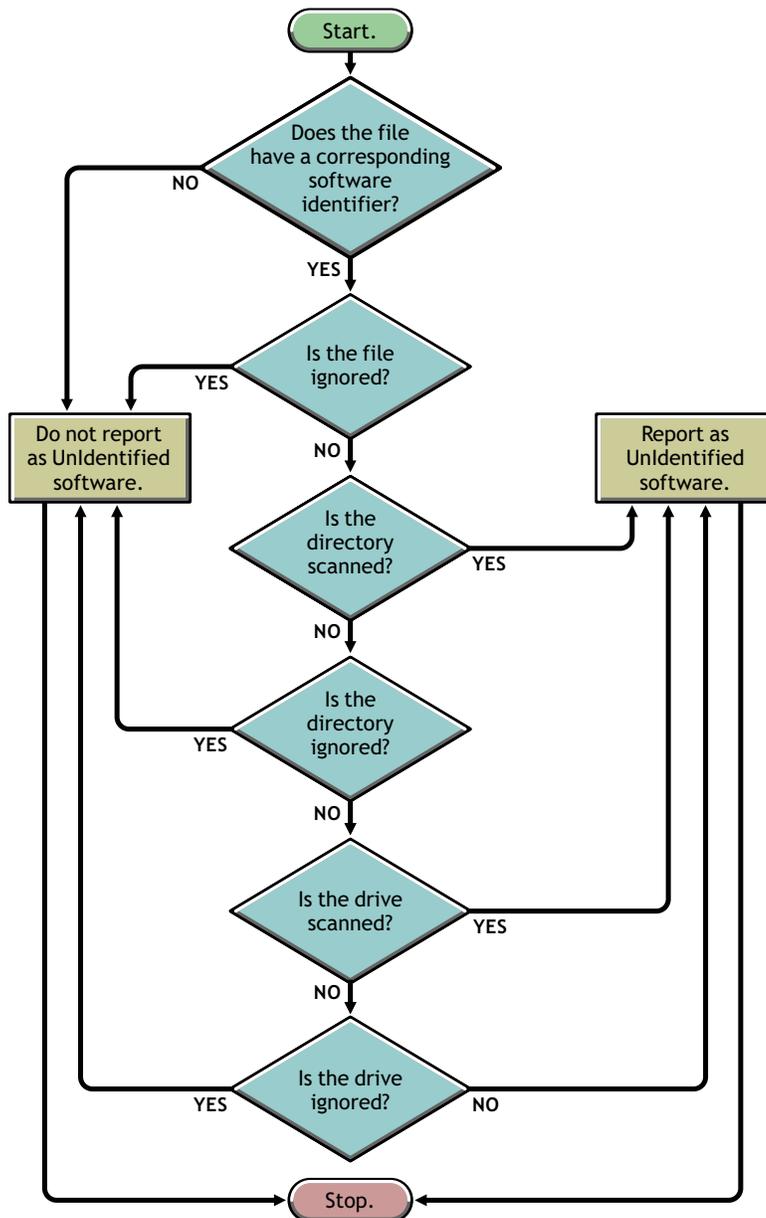


Figure 16-2 Precedence of Software Dictionary rules in the Software Scanning category



Precedence of Software Dictionary Rules in the Disk Usage Scanning Category

The software dictionary rules in the Disk Usage Scanning category determine whether a file should be considered for disk usage scan.

The Disk Usage Scanning category includes the following software dictionary rules:

- ◆ [Section 16.3.16, “Report Disk Space Used by File Extensions,”](#) on page 607
- ◆ [Section 16.3.22, “Disk Usage Scanning Filters - Files,”](#) on page 622
- ◆ [“Scan Directories”](#) on page 621
- ◆ [“Ignore Directories”](#) on page 620

- ◆ “Scan Drives” on page 619
- ◆ “Ignore Drives” on page 619

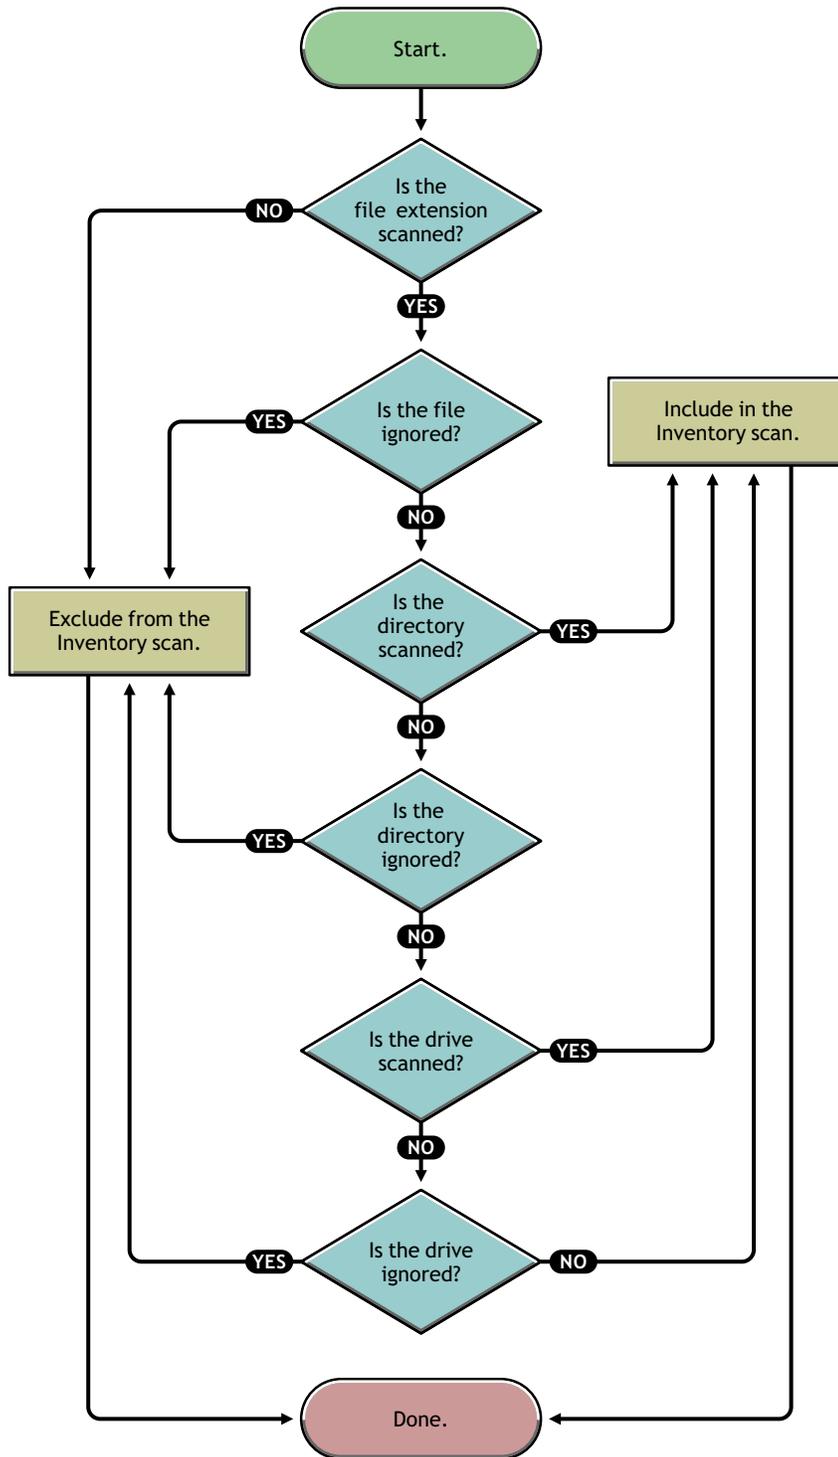
For a file to be considered for the disk usage scan, its file extension must be listed in the “Report Disk Space Used by File Extensions” rule and it should not be excluded from inventory scan in the other Disk Usage Scanning rules.

The following list indicates the precedence of rules in the descending order:

- ◆ Report Disk Space Used by File Extensions
- ◆ Disk Usage Scanning Filters - Files
- ◆ Scan Directories
- ◆ Ignore Directories
- ◆ Scan Drives
- ◆ Ignore Drives

The following flowchart illustrates the precedence of these rules:

Figure 16-3 Precedence of Software Dictionary rules in the Disk Usage Scanning category



16.3.9 Understanding the Software Dictionary Pattern Types

Before configuring the software dictionary rules, you must be aware of the following software dictionary pattern types that are supported in ZENworks 7:

- ♦ “Regular Expression” on page 595
- ♦ “Expandable Expression” on page 595
- ♦ “System Expandable Expression” on page 595

Regular Expression

Regular Expression refers to the POSIX* regular expressions. For more information on regexp (regular expressions), see [The Open Group Base Specifications Issue 6 Web site \(http://www.opengroup.org/onlinepubs/007904975/basedefs/xbd_chap09.html\)](http://www.opengroup.org/onlinepubs/007904975/basedefs/xbd_chap09.html).

Examples of Regular Expression usage:

- ♦ To find all vendor names starting with “Novell,” specify `Novell.*`
- ♦ To find executables, specify `[exe|EXE]`
- ♦ To find files with name containing 6 characters, starting with “r” and ending with “t,” specify `[r...t]`
- ♦ To find files with name starting from A to C, and ending with E, specify `[A-C].*[E]`
- ♦ To find files whose name does not contain any uppercase letters, specify `[^A-Z]+`

NOTE: To use metacharacters such as `[, \, ^, $, ., |, ?, (,), *, and +` as characters, you must prefix them with a backslash (`\`). For example, to specify `c:\windows` as a regular expression, specify it as `c:\\windows`.

Expandable Expression

An Expandable Expression contains displayable characters and the asterisk (*) wildcard character.

“*” matches to zero or more displayable characters.

Examples of Expandable Expression usage:

- ♦ To find all instances of the vendor name beginning with “Microsoft,” specify `Microsoft*`
- ♦ To find files with extension “.exe” in the scan, specify `exe`

System Expandable Expression

- ♦ **On NetWare:** A System Expandable Expression contains displayable characters or references to environmental variables.

Example of an environmental variable: `$sysdir`

- ♦ **On Windows:** A System Expandable Expression contains displayable characters, references to environmental variables, or the asterisk (*) wildcard character.

“*” matches to zero or more displayable characters.

Example of an environmental variable: `%temp%`

IMPORTANT: A System Expandable Expression can contain a combination of displayable characters, references to environmental variables, or the asterisk (*) wildcard character, however, if it contains an environmental variable, you must specify it at the beginning of the expression. For example, %temp%/*

Examples of System Expandable Expression usage:

- ◆ To find the disk usage of the C drive, specify C
- ◆ To find files in the c:\program files directory, specify c:\program files
- ◆ To find files with the extensions, “.com,” specify com

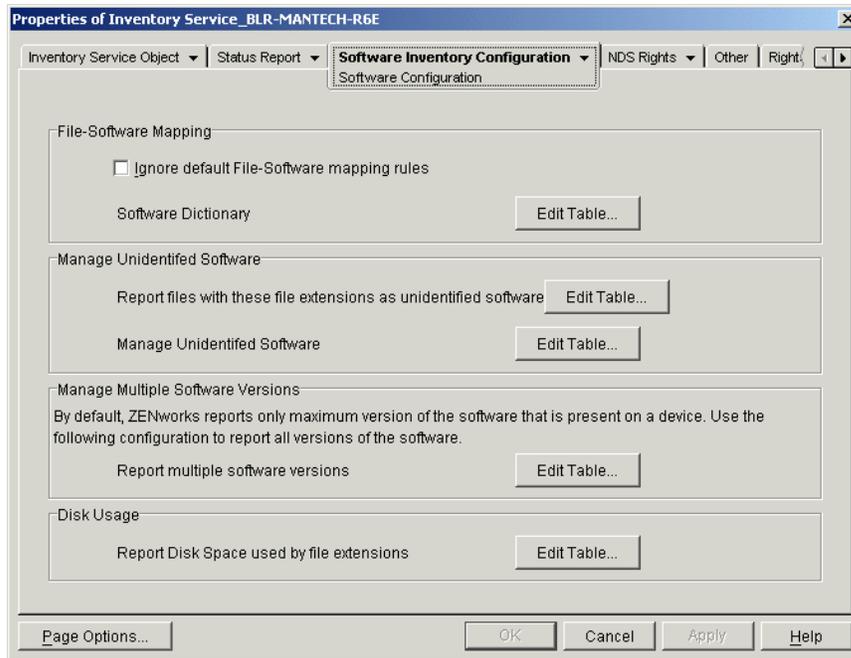
16.3.10 Configuring the Software Dictionary Rules

- 1 In ConsoleOne, right-click the Inventory Service object (Inventory Service_*server_name*), then click *Properties*.
- 2 Click the *Software Inventory Configuration* tab. The Software Configuration page is displayed by default.

You can configure the following settings to scan the software inventory information:

- ◆ **File - Software Mapping:** Includes the following rules:
 - ◆ [Section 16.3.11, “Ignore Default File-Software Mapping Rules,” on page 599](#)
 - ◆ [Section 16.3.12, “Software Dictionary,” on page 599](#)
- ◆ **Manage Unidentified Software:** Includes the following rules:
 - ◆ [Section 16.3.13, “Report Files with These File Extensions As Unidentified Software,” on page 603](#)
 - ◆ [Section 16.3.14, “Manage Unidentified Software,” on page 604](#)
- ◆ **Manage Multiple Software Versions:** Includes the following rule:
 - ◆ [Section 16.3.15, “Report Multiple Software Versions,” on page 605](#)
- ◆ **Disk Usage:** Includes the following rule:
 - ◆ [Section 16.3.16, “Report Disk Space Used by File Extensions,” on page 607](#)

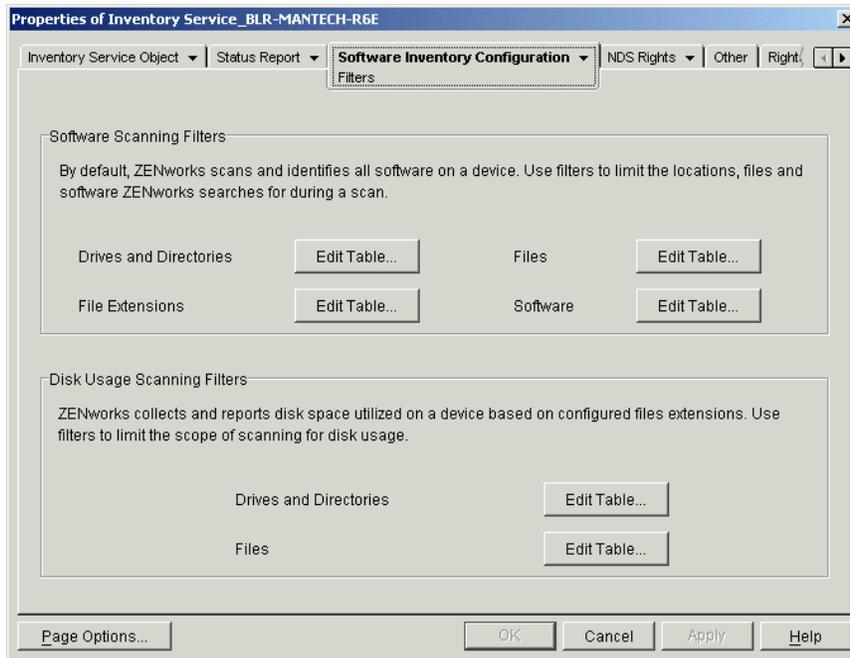
IMPORTANT: Before configuring any ZENworks software dictionary rules, you must be aware of how to use these rules. For detailed information, see [Section 16.3.8, “Understanding the Usage and Precedence of ZENworks Software Dictionary Rules,” on page 589](#).



3 To control the scope of scanning for files, click the *Filters* page and configure the following settings:

- ◆ **Software Scanning Filters:** Includes the following filters:
 - ◆ [Section 16.3.17, “Software Scanning Filters - Drives and Directories,”](#) on page 608
 - ◆ [Section 16.3.18, “Software Scanning Filters - File Extensions,”](#) on page 613
 - ◆ [Section 16.3.19, “Software Scanning Filters - Files,”](#) on page 615
 - ◆ [Section 16.3.20, “Software Scanning Filters - Software,”](#) on page 616
- ◆ **Disk Usage Scanning Filters:** Includes the following filters:
 - ◆ [Section 16.3.21, “Disk Usage Scanning Filters - Drives and Directories,”](#) on page 618
 - ◆ [Section 16.3.22, “Disk Usage Scanning Filters - Files,”](#) on page 622

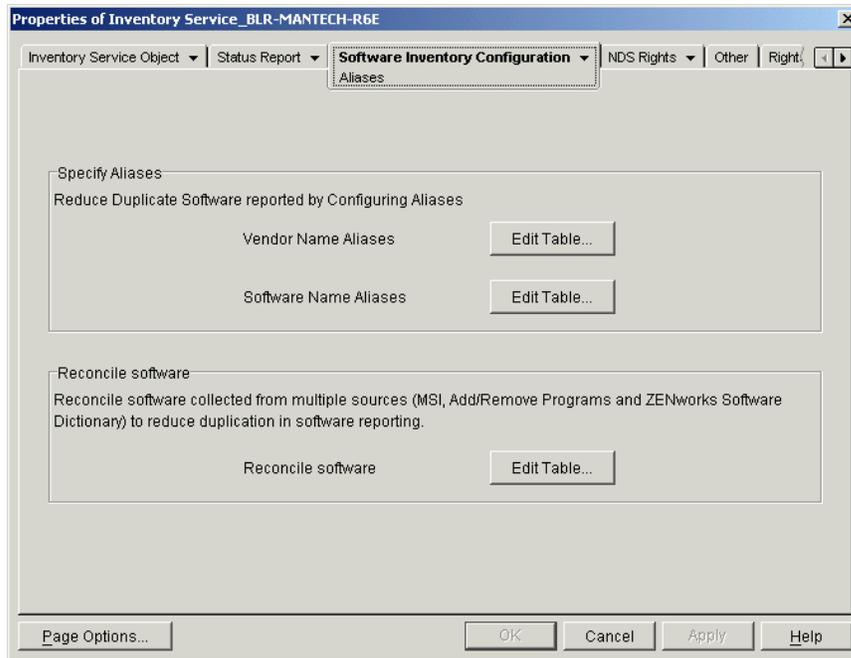
IMPORTANT: Before configuring any ZENworks software dictionary rules, you must be aware of how to use these rules. For detailed information, see [Section 16.3.8, “Understanding the Usage and Precedence of ZENworks Software Dictionary Rules,”](#) on page 589.



4 Click the *Aliases* page to configure the following options:

- ♦ **Specify Aliases:** Allows you to configure aliases for vendor and software names.
By default, the software information is categorized by vendor name in the Inventory ConsoleOne utilities. The software from the same vendor might sometimes have differing vendor names or product names. In this scenario, the Inventory ConsoleOne utilities display the software information under different sections.
However, you can merge the software information by specifying aliases. You customize these settings in the following software dictionary rules:
 - ♦ [Section 16.3.23, “Vendor Name Aliases,” on page 623](#)
 - ♦ [Section 16.3.24, “Software Name Aliases,” on page 625](#)
- ♦ **Reconcile Software:** Allows you to merge the software identified through Add/Remove Programs or the MSI, and the software identified through the ZENworks software dictionary. For more information, see [Section 16.3.25, “Reconcile Software,” on page 626](#).

IMPORTANT: Before configuring any ZENworks software dictionary rules, you must be aware of how to use these rules. For detailed information, see [Section 16.3.8, “Understanding the Usage and Precedence of ZENworks Software Dictionary Rules,” on page 589](#).



5 Click *Apply*, then click *Close*.

16.3.11 Ignore Default File-Software Mapping Rules

In the *Software Configuration* property page, select the *Ignore Default File-Software Mapping Rules* check box if you do not want the Inventory scanner to use the default File-Software mapping rules that are configured in the ZENworks software dictionary for scanning software inventory information.

IMPORTANT: This option is not available for selection if the software dictionary is updated from another Inventory server.

16.3.12 Software Dictionary

The *Software Dictionary* option allows you to configure software identifiers in the ZENworks software dictionary.

By default, the ZENworks software dictionary contains predefined software identifiers. You can create new software identifiers in the ZENworks software dictionary by editing the predefined software identifiers or creating a new software identifier.

To configure rules in the ZENworks software dictionary:

- 1 In the *Software Configuration* property page, Click the *Edit Table* option of *Software Dictionary*.

The Software Dictionary table is displayed.

| Dictionary ... | Key identifier | Filename | File Last Modified Time... | Minimum file size | Maximu... | Softw: |
|----------------|----------------|-------------|----------------------------|-------------------|-----------|------------|
| 33823 | No | MVREAD... | 2003-07-14 15:25 | 2460160 | 2460160 | Reader |
| 33822 | No | NETSONI... | 2000-12-18 12:56 | 3000096 | 3000096 | Netsonic |
| 33821 | No | MSHOW... | 2003-10-29 11:37 | 639056 | 639056 | Mshow |
| 33820 | No | PCBODY... | 2003-01-20 18:06 | 942080 | 942080 | PC Body |
| 33819 | No | AGMAIL.E... | 2004-01-19 23:29 | 1511424 | 1511424 | Group M |
| 33818 | No | NOTETA... | 2002-08-26 17:26 | 1725440 | 1725440 | NoteTab |
| 33817 | No | CITYDES... | 2002-08-08 17:21 | 3891200 | 3891200 | CityDesk |
| 33816 | No | EEBED9... | 2003-08-17 22:22 | 294912 | 294912 | Executab |
| 33815 | No | EXEAPI1... | 2003-08-17 22:31 | 17120 | 17120 | Executab |
| 33814 | No | PIM.EXE | 2002-12-02 19:01 | 1646592 | 1646592 | Mp3 Play |
| 33813 | No | 2020.EXE | 2001-05-08 16:21 | 2085376 | 2085376 | 20/20 |
| 33812 | No | BSEERAS... | 2003-09-22 11:18 | 208896 | 208896 | Erase |
| 33811 | No | ACU.EXE | 2003-09-26 11:29 | 1339392 | 1339392 | ACU |
| 33810 | No | PCARMD... | 2002-04-18 01:17 | 45056 | 45056 | AMBIT WI |
| 33809 | No | MAINCTR... | 2003-08-06 08:32 | 327680 | 327680 | Silence Ir |
| 33808 | No | KILLAD.E... | 2000-01-27 16:00 | 30720 | 30720 | KillAd |
| 33807 | No | IPHOTON... | 2003-11-30 23:00 | 1458176 | 1458176 | Newsgro |
| 33806 | No | IPHOTO... | 2003-10-28 22:50 | 1839104 | 1839104 | lphoto |

Buttons: Insert, Delete, Sort..., Filter..., OK, Cancel, Help

Right Panel: Add to (Ignore Software), Add From (Unidentified Software)

The Software Dictionary table displays the data stored in the ZENworks software dictionary. It might contain entries that are:

- ♦ **Light gray in color:** Indicates that these entries will not be considered in a scan because the table already contains entries that override these entries.
- ♦ **Dark gray in color:** These are inherited rules. For more information about inherited rules, see [Section 16.3.6, “What is an Inherited Rule?”](#) on page 588.

2 In the Software Dictionary table, you can perform the following operations:

- ♦ [“Manually Adding Entries to the Software Dictionary”](#) on page 600
- ♦ [“Automatically Adding Entries to the Software Dictionary”](#) on page 602
- ♦ [“Deleting Entries from the Software Dictionary”](#) on page 602
- ♦ [“Modifying the Values of the Software Dictionary Entries”](#) on page 602
- ♦ [“Excluding a Software from a Scan”](#) on page 603
- ♦ [Section 16.3.26, “Sorting Entries in the Table,”](#) on page 627
- ♦ [Section 16.3.27, “Filtering Entries in the Table,”](#) on page 627
- ♦ [Section 16.3.28, “Refreshing Entries in the Table,”](#) on page 628

3 Click *OK*.

Manually Adding Entries to the Software Dictionary

1 In the Software Dictionary table, click *Insert* to add a new row.

2 Specify values for the following attributes:

Filename, File Last Modified Time (yyyy-dd-mm hours:minutes), Minimum File Size (bytes), Maximum File Size (bytes), Software Name, Support Pack, Software Version, Internal Version, Description, Vendor, Platform, and Category.

The following attributes are called “matching attributes”: Filename, File Last Modified Time, Minimum File Size, Maximum File Size, and Internal Version. The values of these matching attributes are compared with the values scanned by the Inventory scanner from the file headers on the inventoried servers. If the values are same, the values in the corresponding software

information attributes (Software Name, Support Pack, Software Version, Description, Vendor, Platform, and Category) are stored in the Inventory database.

In the Software Dictionary table, you must specify values for the following attributes: Filename, Software Name, and Vendor. It is optional to specify values for other attributes.

When you add an entry, a unique ID called the Dictionary Identifier is automatically assigned to this entry.

For example, configure the following settings in the Software Dictionary - Row Editor table:

Filename= MSACCESS . EXE
File Last Modified Time = 1998-30-01 05:30
Minimum File Size = 299854
Maximum File Size = 400000
Software Name = Access
Software Version = 7.0
Internal Version = 7.0
Description = Microsoft Access
Vendor = Microsoft
Category = Database

If the Inventory scanner finds a file with the following values during the scan: "File Name= MSACCESS . EXE; File Last Modified Time= 1998-30-01 05:30; File Size= 300000", then the following information is stored in the Inventory database:

Software Name = Access
Software Version = 7.0
Description = Microsoft Access
Vendor = Microsoft
Category = Database

If you do not specify a value for an attribute, then this attribute is not considered to determine the overriding entry. Also, only the matching attributes are considered to determine the overriding entry. For example, the Configure Dictionary table has the following entries for MS Word:

| Filename | Minimum File Size | Maximum File Size | Software Name | Vendor |
|-------------|-------------------|-------------------|---------------|-----------|
| winword.exe | 10000 | 10000 | Word | Microsoft |
| winword.exe | 0 | 30000 | Word | Microsoft |

To determine the overriding entry, only the maximum file size value is considered. Consequently, the second entry with 30000 maximum file size overrides the first entry.

- 3 (Optional) Select the *Key Identifier* check box for this entry.

For example, the Software Dictionary table has the following entries for MS Word:

| Filename | File Last Modified Time | Minimum File Size | Maximum File Size | Software Name | Software version | Internal version | Description | Vendor |
|-------------|-------------------------|-------------------|-------------------|---------------|------------------|------------------|-------------------------------|-----------|
| winword.exe | 2004-30-10 5:30 | 10000 | 10000 | Word | 2002 | 10.0.4219 | Microsoft Word | Microsoft |
| osa.exe | 2004-30-02 16:00 | 10000 | 10000 | Word | 2002 | 10.0.4300 | Microsoft Office XP Component | Microsoft |

If the key identifier has not been defined, the software information for MS Word might be selected from anyone of the above entries.

To ensure that the information from the identifier corresponding to “Winword.exe” is selected, select Key Identifier for “Winword.exe.” If you select “Winword.exe” as the key identifier in the Configure Software Dictionary table, the Inventory scanner stores the information related to Winword.exe into the Inventory database.

Automatically Adding Entries to the Software Dictionary

- 1 In the Software Dictionary table, click *Unidentified Software* located in the *Add From* pane.
- 2 In the Manage Unidentified Software table, do the following:
 - 2a Select the entry to be added to the software dictionary.
 - 2b Click *Software Dictionary* located in the *Add To* pane.
 - 2c Click *Close*.

Deleting Entries from the Software Dictionary

- 1 In the Software Dictionary table, select the entry to be deleted.
- 2 Click *Delete*.

IMPORTANT: You can delete only the non-inherited entries.

Modifying the Values of the Software Dictionary Entries

- 1 In the Software Dictionary table, double-click the entry whose values you want to modify.
You can modify only one entry at a time.

TIP: You can also invoke the Row Editor dialog box by selecting the entry you want to modify and pressing either one of the keys: Enter, Spacebar, or F2.

- 2 Modify the values.
You cannot modify the values of the Dictionary Identifier and Filename attributes.
- 3 Click *OK*.

IMPORTANT: You cannot modify the values of an inherited rule. Also, modifying a default predefined rule creates a new user-defined rule.

Excluding a Software from a Scan

- 1 In the Software Dictionary table, select the corresponding entry for the software you want to exclude from the Inventory scan.
- 2 Click *Ignore Software* located in the *Add To* pane.

The entry is added to the Ignore Software table in [Software Scanning Filters - Software](#).

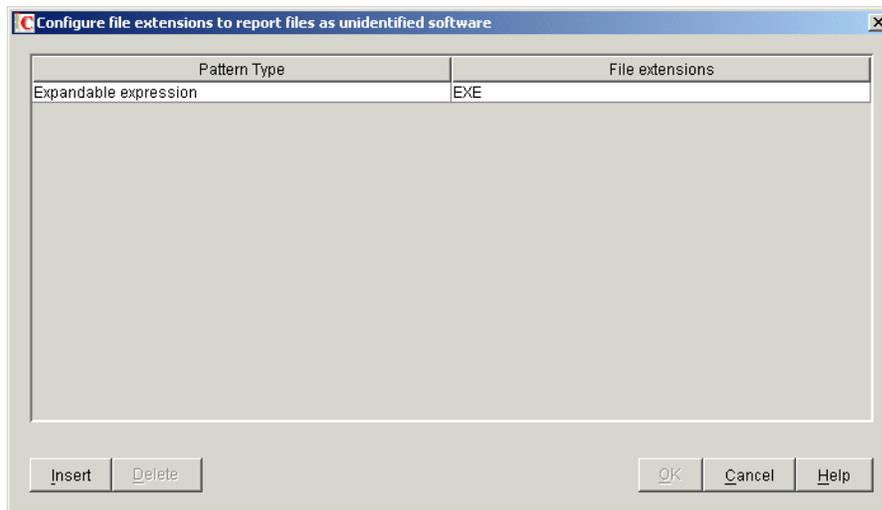
16.3.13 Report Files with These File Extensions As Unidentified Software

The “Report Files with These File Extensions As Unidentified Software” rule allows you to configure file extension of files that must be reported as unidentified software.

To configure the rule:

- 1 In the *Software Configuration* property page, click the *Edit Table* option of *Report Files with These File Extensions As Unidentified Software*.

The “Configure File Extensions to Report Files as Unidentified Software” table is displayed.



- 2 Click *Insert* to add a new row.
- 3 In the *Pattern Type* drop-down list, select *Expandable Expression* or *Regular Expression*.
- 4 Specify a file extension.
- 5 Click *OK*.

For example, if you want the Inventory scanner to report the software with the `.exe` extension as Unidentified software, configure the following settings in the table:

Pattern Type = Expandable Expression

File Extensions = exe

IMPORTANT: The table might contain **inherited rules**. You can edit or delete these rules only in the software dictionary of the inventory server configured in the Dictionary Update policy. These

rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

To delete an entry from the table, select the entry and click Delete. You can delete only the non-inherited entries.

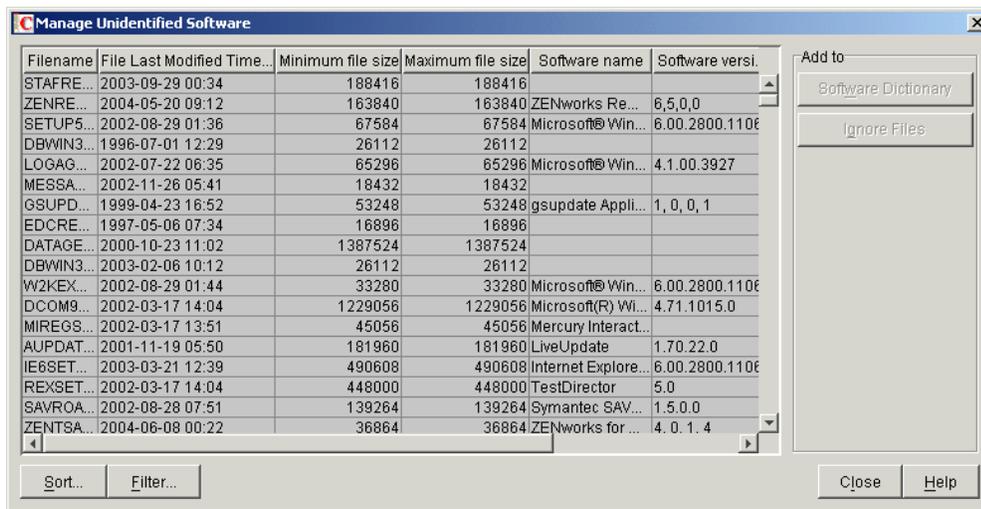
16.3.14 Manage Unidentified Software

The “Manage Unidentified Software” rule allows you to include or exclude the **unidentified software** from the inventory scan.

To configure this rule:

- 1 In the *Software Configuration* property page, click the *Edit Table* option of *Manage Unidentified Software*.

The Manage Unidentified Software table is displayed.



| Filename | File Last Modified Time... | Minimum file size | Maximum file size | Software name | Software versi... |
|-----------|----------------------------|-------------------|-------------------|---------------------|-------------------|
| STAFRE... | 2003-09-29 00:34 | 188416 | 188416 | | |
| ZENRE... | 2004-05-20 09:12 | 163840 | 163840 | ZENworks Re... | 6.5.0.0 |
| SETUP5... | 2002-08-29 01:36 | 67584 | 67584 | Microsoft® Win... | 6.00.2800.1106 |
| DBWIN3... | 1996-07-01 12:29 | 26112 | 26112 | | |
| LOGAG... | 2002-07-22 06:35 | 65296 | 65296 | Microsoft® Win... | 4.1.00.3927 |
| MESSA... | 2002-11-26 05:41 | 18432 | 18432 | | |
| GSUPD... | 1999-04-23 16:52 | 53248 | 53248 | gsupdate Appli... | 1, 0, 0, 1 |
| EDCRE... | 1997-05-06 07:34 | 16896 | 16896 | | |
| DATAGE... | 2000-10-23 11:02 | 1387524 | 1387524 | | |
| DBWIN3... | 2003-02-06 10:12 | 26112 | 26112 | | |
| W2KEX... | 2002-08-29 01:44 | 33280 | 33280 | Microsoft® Win... | 6.00.2800.1106 |
| DCOM9... | 2002-03-17 14:04 | 1229056 | 1229056 | Microsoft(R) Wi... | 4.71.1015.0 |
| MIREGS... | 2002-03-17 13:51 | 45056 | 45056 | Mercury Interact... | |
| AUPDAT... | 2001-11-19 05:50 | 181960 | 181960 | LiveUpdate | 1.70.22.0 |
| IE6SET... | 2003-03-21 12:39 | 490608 | 490608 | Internet Explore... | 6.00.2800.1106 |
| REXSET... | 2002-03-17 14:04 | 448000 | 448000 | TestDirector | 5.0 |
| SAVROA... | 2002-08-28 07:51 | 139264 | 139264 | Symantec SAV... | 1.5.0.0 |
| ZENTSA... | 2004-06-08 00:22 | 36864 | 36864 | ZENworks for ... | 4.0.1.4 |

IMPORTANT: The table might contain **inherited rules**. You can edit or delete these rules only in the software dictionary of the inventory server configured in the Dictionary Update policy. These rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

- 2 In the Manage Unidentified Software table, you can perform the following operations:
 - ♦ “Including Unidentified Software in a Scan” on page 605
 - ♦ “Excluding Unidentified Software from the Scan” on page 605
 - ♦ Section 16.3.26, “Sorting Entries in the Table,” on page 627
 - ♦ Section 16.3.27, “Filtering Entries in the Table,” on page 627
 - ♦ Section 16.3.28, “Refreshing Entries in the Table,” on page 628
- 3 Click *OK*.

Including Unidentified Software in a Scan

If you want unidentified software to be reported as a known software in subsequent scans, do the following:

- 1 Select the software entry in the Manage Unidentified Software table.
- 2 Click *Software Dictionary* located in the *Add To* pane.

The entry is automatically added to the **Software Dictionary** table.

Excluding Unidentified Software from the Scan

If you want unidentified software not to be reported in subsequent scans, do the following:

- 1 Select the software entry in the Manage Unidentified Software table.
- 2 Click *Ignore Files* located in the *Add To* pane.

The entry is automatically added to the table in **Software Scanning Filters - Files**.

16.3.15 Report Multiple Software Versions

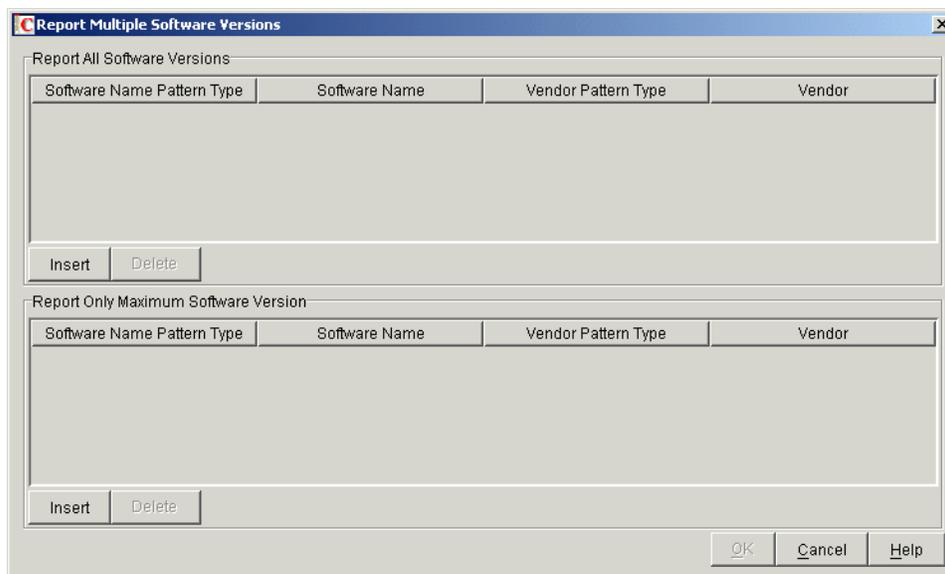
The “Report Multiple Software Versions” rule allows you to specify the software for which the Inventory scanner must report multiple versions installed on the inventoried server.

By default, the Inventory scanner scans for the highest version of the software installed on the inventoried server.

To configure this rule:

- 1 In the *Software Configuration* property page, click the *Edit Table* option of *Report Multiple Software Versions*.

The Report Multiple Software Versions dialog box is displayed.



- 2 If you want the Inventory scanner to report all versions of the software installed on the inventoried servers, configure a rule in the Report All Software Versions table.

- 2a** In the Report All Software Versions table, click *Insert* to add a new row.
- 2b** In the *Software Name Pattern Type* drop-down list, select *Expandable Expression* or *Regular Expression*.
- 2c** Specify a software name.
- 2d** (Optional) In the *Vendor Pattern Type* drop-down list, select *Expandable Expression* or *Regular Expression*.
- 2e** (Optional) Specify a vendor name.

For example, if you want the Inventory scanner to report all versions of the Adobe Acrobat Reader installed on the inventoried server, configure the following settings in the table:

Software Name Pattern Type = Expandable Expression
Software Name = Acrobat* Reader*
Vendor Pattern Type = Expandable Expression
Vendor Name = Adobe*

If the inventoried server has Acrobat Reader versions 5.0 and 6.0 installed, the Inventory scanner reports both versions of Acrobat Reader (5.0 and 6.0).

IMPORTANT: The table might contain **inherited rules**. You can edit or delete these rules only in the software dictionary of the inventory server configured in the Dictionary Update policy. These rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

To delete an entry from the table, select the entry and click *Delete*. You can delete only the non-inherited entries.

- 3** If you want the Inventory scanner to report only the highest version of the software installed on the inventoried servers, configure a rule in the Report Only Maximum Software Version table.
 - 3a** In the Report Only Maximum Software Version table, click *Insert* to add a new row.
 - 3b** In the *Software Name Pattern Type* drop-down list, select *Expandable Expression* or *Regular Expression*.
 - 3c** Specify a software name.
 - 3d** (Optional) In the *Vendor Pattern Type* drop-down list, select *Expandable Expression* or *Regular Expression*.
 - 3e** (Optional) Specify a vendor name.

For example, if you want the Inventory scanner to report only the highest version of the Adobe Acrobat Reader installed on the inventoried server, configure the following settings in the table:

Software Name Pattern Type = Expandable Expression
Software Name = Acrobat* Reader*
Vendor Pattern Type = Expandable Expression
Vendor Name= Adobe*

If the inventoried server has Adobe Acrobat Reader versions 4.0 and 5.0 installed, then the Inventory scanner reports only Adobe Acrobat Reader 5.0.

IMPORTANT: The table might contain **inherited rules**. You can edit or delete these rules only in the software dictionary of the inventory server configured in the Dictionary

Update policy. These rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

To delete an entry from the table, select the entry and click *Delete*. You can delete only the non-inherited entries.

4 Click *OK*.

By default, the scanner reports only the highest version of the software installed. If a rule in Report All Software Versions conflicts with a rule in Report Only Maximum Software Version, then the rule in Report Only Maximum Software Version overrides the rule of Report All Software Versions.

For example, if you want the Inventory scanner to report all versions of Microsoft software except for Microsoft Office, and also report only the highest version of Microsoft Office installed, configure the following filters as shown below:

- ◆ **Report All Software Versions:** Configure the following settings:

Software Name Pattern Type = Expandable Expression

Software Name = *

Vendor Pattern Type = Expandable Expression

Vendor Name= Microsoft*

- ◆ **Report Only Maximum Version:** Configure the following settings:

Software Name Pattern Type = Expandable Expression

Software Name = *office*

Vendor Pattern Type = Expandable Expression

Vendor Name= Microsoft*

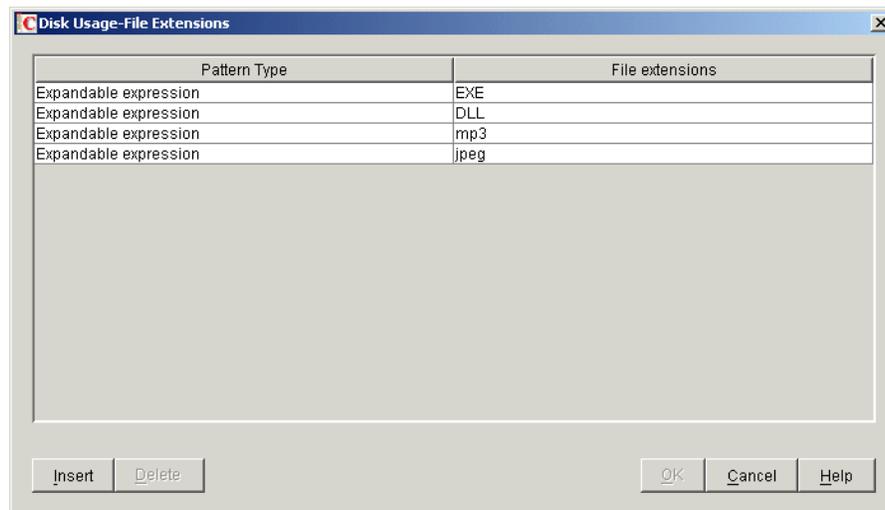
16.3.16 Report Disk Space Used by File Extensions

The “Report Disk Space Used by File Extensions” rule allows you to specify the file extension of the files whose total disk usage you want to scan.

To configure this rule:

- 1 In the *Software Configuration* property page, click the *Edit Table* option of *Report Disk Space Used by File Extensions*.

The Disk Usage - File Extensions table is displayed.



- 2 Click *Insert* to add a new row.
- 3 In the *Pattern Type* drop-down list, select *Expandable Expression* or *Regular Expression*.
- 4 Specify a file extension.
- 5 Click *OK*.

For example, if you want the Inventory scanner to scan for disk usage of all files with extension `.pif`, configure the following settings in the Disk Usage - File Extensions table:

Pattern Type = Expandable Expression

File Extension = pif

The Inventory scanner scans and stores only the total disk usage for all files with extension `.pif` in the Inventory database.

IMPORTANT: The table might contain **inherited rules**. You can edit or delete these rules only in the software dictionary of the inventory server configured in the Dictionary Update policy. These rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

To delete an entry from the Disk Usage - File Extensions table, select the entry and click *Delete*. You can delete only the non-inherited entries.

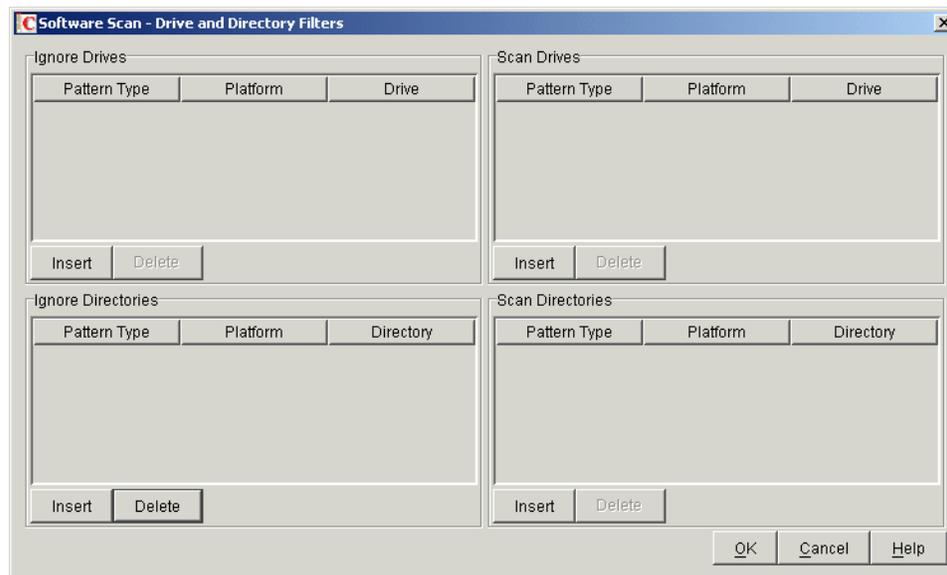
16.3.17 Software Scanning Filters - Drives and Directories

The “Drives and Directories” filter the allows you to control the scanning for software files located in specified drives and directories.

To configure this filter:

- 1 In the *Filters* property page, click the *Edit Table* option of *Drives and Directories* located in the *Software Scanning Filters* pane.

The Software Scan - Drive and Directory Filters dialog box is displayed.



2 Configure the following filters:

- ◆ "Ignore Drives" on page 610
- ◆ "Scan Drives" on page 610
- ◆ "Ignore Directories" on page 611
- ◆ "Scan Directories" on page 612

By default, the Inventory scanner scans all directories on the inventoried servers. If you have configured a rule that ignores all directories during a scan by using the Ignore Directories filter, but now want to include a specific directory in a scan, you can identify the specific directory using the Scan Directories filter. The settings of the Scan Directories filter overrides the settings of the Ignore Directories and Ignore Drives filters.

For example, if you want the Inventory scanner to ignore all files and directories in C : except for the c:\program files directory on Windows inventoried servers, configure the following filters as shown below:

- ◆ **Ignore Drives:** Configure the following settings:
 - Pattern Type = System Expandable Expression
 - Platform = Windows
 - Drive = C
- ◆ **Scan Directories:** Configure the following settings:
 - Pattern Type = System Expandable Expression
 - Platform = Windows
 - Drive = c:\program files

3 Click *OK*.

Ignore Drives

The “Ignore Drives” filter allows you to specify the drives that should not be scanned for on the inventoried servers.

By default, the Inventory scanner scans all drives.

To configure this filter:

- 1 In the Ignore Drives table, click *Insert* to add a new row.
- 2 In the *Pattern Type* drop-down list, select *System Expandable Expression*, *Expandable Expression*, or *Regular Expression*.
- 3 (Conditional) If you select *System Expandable Expression* as the pattern type, then select *NetWare* or *Windows* in the *Platform* drop-down list, depending on the operating system of the inventoried servers.

IMPORTANT: If you select *Expandable Expression* or *Regular Expression* as the pattern type, the corresponding value in the *Platform* column automatically changes to *Any*. You cannot change the value.

- 4 Specify a drive name.

For example, if you want the Inventory scanner not to scan the files in C drive on all the Windows inventoried servers, configure the following settings in the Ignore Drives table:

Pattern Type = System Expandable Expression

Platform = Windows

Drive = C

The Inventory scanner does not scan the files in the C drive.

IMPORTANT: The table might contain **inherited rules**. You can edit or delete these rules only in the software dictionary of the inventory server configured in the Dictionary Update policy. These rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

To delete an entry from the Ignore Drives table, select the entry and click *Delete*. You can delete only the non-inherited entries.

Scan Drives

The “Scan Drives” filter allows you to specify the drives that should be scanned for at the inventoried servers.

To configure this filter:

- 1 In the Scan Drives table, click *Insert* to add a new row.
- 2 In the *Pattern Type* drop-down list, select *System Expandable Expression*, *Expandable Expression*, or *Regular Expression*.
- 3 (Conditional) If you select *System Expandable Expression* as the pattern type, then select *NetWare* or *Windows* in the *Platform* drop-down list, depending on the operating system of the inventoried servers.

IMPORTANT: If you select *Expandable Expression* or *Regular Expression* as the pattern type, the corresponding value in the *Platform* column automatically changes to *Any*. You cannot change the value.

4 Specify a drive name.

For example, if you want the Inventory scanner to scan for files in the C drive on all the Windows inventoried servers, configure the following settings in the Scan Drives table:

Pattern Type = System Expandable Expression

Platform = Windows

Drive = C

You must also configure the following settings in the Ignore Drives table:

Pattern Type = System Expandable Expression

Platform = Windows

Drive = *

The Inventory scanner scans only the files in the C drive for the software information.

IMPORTANT: The table might contain **inherited rules**. You can edit or delete these rules only in the software dictionary of the inventory server configured in the Dictionary Update policy. These rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

To delete an entry from the Scan Drives table, select the entry and click *Delete*. You can delete only the non-inherited entries.

IMPORTANT: By default, the Inventory scanner scans all drives on the inventoried servers. If you have configured all drives to be ignored during a scan by using the Ignore Drives filter, but now want to include a specific drive in a scan, you can identify the specific drive using the Scan Drives filter. The settings of the Scan Drives filter override the settings of the Ignore Drives filter.

Ignore Directories

The “Ignore Directories” filter allows you to specify the directories that should not be scanned for at the inventoried servers.

By default, the Inventory scanner scans all directories.

To configure this filter:

- 1 In the Ignore Directories table, click *Insert* to add a new row.
- 2 In the *Pattern Type* drop-down list, select *System Expandable Expression*, *Expandable Expression*, or *Regular Expression*.
- 3 (Conditional) If you select *System Expandable Expression* as the pattern type, then select *NetWare* or *Windows* in the *Platform* drop-down list, depending on the operating system of the inventoried servers.

IMPORTANT: If you select *Expandable Expression* or *Regular Expression* as the pattern type, the corresponding value in the *Platform* column automatically changes to *Any*. You cannot change the value.

4 Specify a directory name.

For example, if you do not want the Inventory scanner to scan the files in the `c:\program files` directory on all the Windows inventoried servers, configure the following settings in the Ignore Directories table:

Pattern Type = System Expandable Expression

Platform = Windows

Directory= C:\Program Files

The Inventory scanner does not scan for the files in `c:\program files`.

IMPORTANT: The table might contain **inherited rules**. You can edit or delete these rules only in the software dictionary of the inventory server configured in the Dictionary Update policy. These rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

To delete an entry from the Ignore Directories table, select the entry and click *Delete*. You can delete only the non-inherited entries.

Scan Directories

The “Scan Directories” filter allows you to specify the directories that should be scanned for at the inventoried servers.

To configure this filter:

- 1** In the Scan Directories table, click *Insert* to add a new row.
- 2** In the *Pattern Type* drop-down list, select *System Expandable Expression*, *Expandable Expression*, or *Regular Expression*.
- 3** (Conditional) If you select *System Expandable Expression* as the pattern type, then select *NetWare* or *Windows* in the *Platform* drop-down list, depending on the operating system of the inventoried servers.

IMPORTANT: If you select *Expandable Expression* or *Regular Expression* as the pattern type, the corresponding value in the *Platform* column automatically changes to *Any*. You cannot change the value.

4 Specify a directory name.

For example, if you want the Inventory scanner to scan for files in the `c:\program files` directory on all the Windows inventoried servers, configure the following settings in the Scan Directories table:

Pattern Type = System Expandable Expression

Platform = Windows

Directory = C:\Program Files

You must also configure the following settings in the Ignore Directories table:

Pattern Type = System Expandable Expression

Platform = Windows

Directory = *

The Inventory scanner scans only the files in `c:\program files` for software information.

IMPORTANT: The table might contain **inherited rules**. You can edit or delete these rules only in the software dictionary of the inventory server configured in the Dictionary Update policy. These rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

To delete an entry from the Scan Directories table, select the entry and click *Delete*. You can delete only the non-inherited entries.

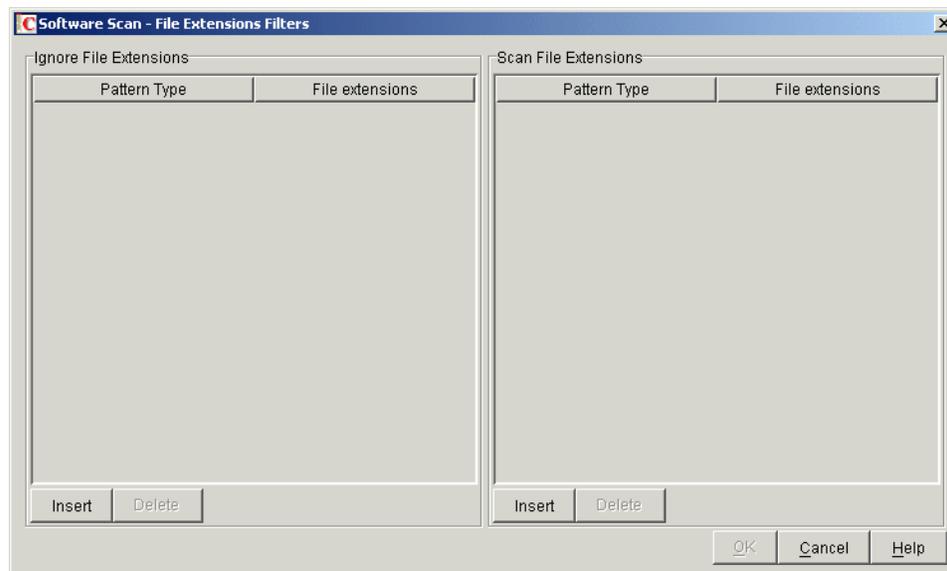
16.3.18 Software Scanning Filters - File Extensions

The “File Extensions” filter allows you to control scanning for software files with a specified extension.

To configure this filter:

- 1 In the *Filters* property page, click the *Edit Table* option of *File Extensions* located in the *Software Scanning Filters* pane.

The Software Scan - File Extensions Filters dialog box is displayed.



- 2 Configure the following filters:
 - ♦ “Ignore File Extensions” on page 614
 - ♦ “Scan File Extensions” on page 614
- 3 Click *OK*.

Ignore File Extensions

The “Ignore File Extensions” filter allows you to specify the file extensions that should not be scanned for at the inventoried servers.

To configure this filter:

- 1 In the Ignore File Extensions table, click *Insert* to add a new row.
- 2 In the *Pattern Type* drop-down list, select *Expandable Expression* or *Regular Expression*.
- 3 Specify a file extension.

For example, if you do not want the Inventory scanner to scan for files whose extension begins with `.ex`, configure the following settings in the Ignore File Extensions table:

Pattern Type = Expandable Expression

File Extension = `ex*`

The Inventory scanner does not scan for the files whose extension begin with `.ex`. For example, `.ex1`, `.ex2`, `.exe`, and `exec`.

IMPORTANT: The table might contain **inherited rules**. You can edit or delete these rules only in the software dictionary of the inventory server configured in the Dictionary Update policy. These rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

To delete an entry from the Ignore File Extensions table, select the entry and click *Delete*. You can delete only the non-inherited entries.

Scan File Extensions

The “Scan File Extensions” filter allows you to specify the file extensions that should be scanned for at the inventoried servers.

If you have excluded file extensions from scanning by using the Ignore File Extensions filter, but now want to include a specific file extension in the scan, you can identify the specific file extension using the Scan File Extensions filter. The settings of the Scan File Extensions filter override the settings of the Ignore File Extensions filter.

To configure this filter:

- 1 In the Scan File Extensions table, click *Insert* to add a new row.
- 2 In the *Pattern Type* drop-down list, select *Expandable Expression* or *Regular Expression*.
- 3 Specify a file extension.

For example, if you want the Inventory scanner to scan for all files with a `.exe` extension, configure the following settings in the Scan File Extension table:

Pattern Type = Regular Expression

File Extension = `[exe|EXE]`

The Inventory scanner scans and stores only the files with extension `.exe` in the Inventory database.

IMPORTANT: The table might contain **inherited rules**. You can edit or delete these rules only in the software dictionary of the inventory server configured in the Dictionary Update policy. These rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

To delete an entry from the Scan File Extensions table, select the entry and click *Delete*. You can delete only the non-inherited entries.

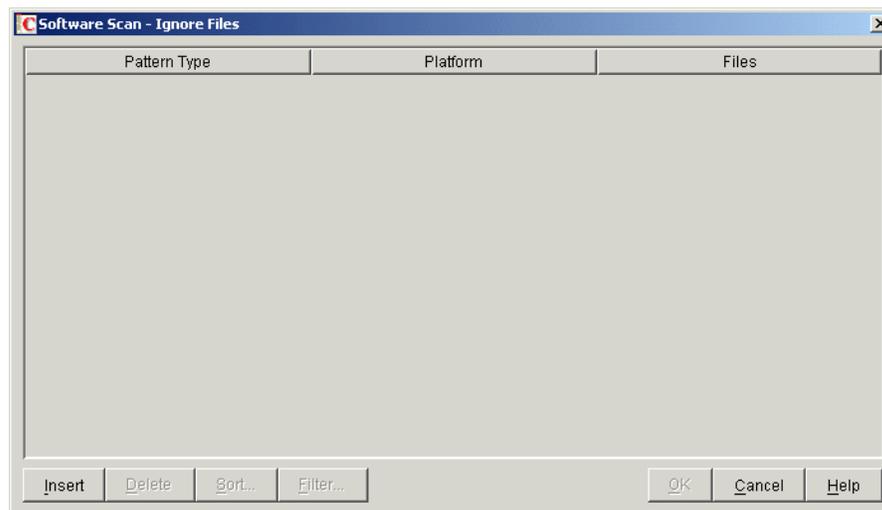
16.3.19 Software Scanning Filters - Files

The “Files” filter allows you to configure software to be excluded during the Inventory scan.

To configure this filter:

- 1 In the *Filters* property page, click the *Edit Table* option of *Files* located in the *Software Scanning Filters* pane.

The Software Scan - Ignore Files table is displayed.



- 2 Click *Insert* to add a new row.
- 3 In the *Pattern Type* drop-down list, select *System Expandable Expression*, *Expandable Expression*, or *Regular Expression*.
- 4 (Conditional) If you select *System Expandable Expression* as the pattern type, then select *NetWare* or *Windows* in the *Platform* drop-down list, depending on the operating system of the inventoried servers.

IMPORTANT: If you select *Expandable Expression* or *Regular Expression* as the pattern type, the corresponding value in the *Platform* column is automatically changed to *Any*. You cannot change the value.

- 5 Specify a filename.
- 6 Click *OK*.

For example, if you want the Inventory scanner to scan `notepad.exe` on all the Windows inventoried servers, configure the following settings:

Platform = Windows
Pattern Type = System Expandable Expression
Files = notepad.exe

This table also displays files that are added from the Manage Unidentified Software table.

IMPORTANT: The table might contain **inherited rules**. You can edit or delete these rules only in the software dictionary of the inventory server configured in the Dictionary Update policy. These rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

You can also perform the following operations in the Software Scan - File Filters table:

- ♦ Deleting only the non-inherited entries.
- ♦ **Sorting Entries in the Table.**
- ♦ **Filtering Entries in the Table.**
- ♦ **Refreshing Entries in the Table.**

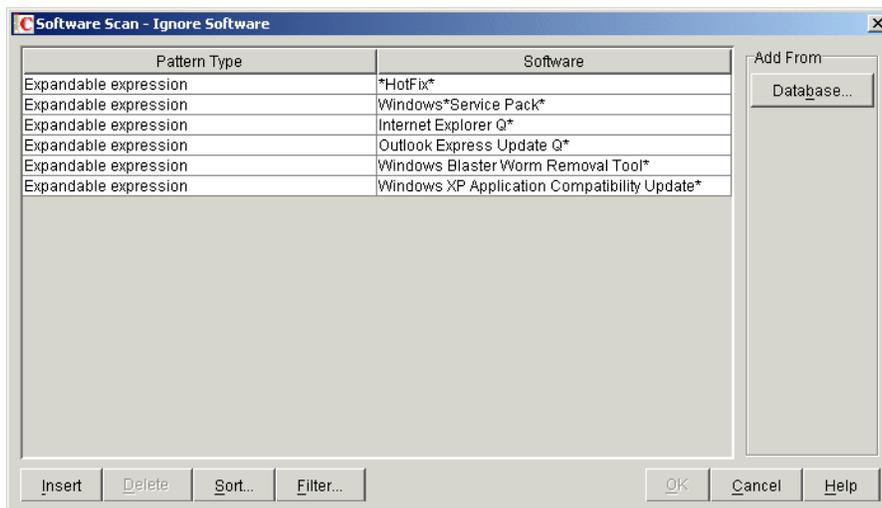
16.3.20 Software Scanning Filters - Software

The “Software” filter allows you to configure a software that is not to be reported during the Inventory scan.

To configure this filter:

- 1 In the *Filters* property page, click the *Edit Table* option of *Software* located in the *Software Scanning Filters* pane.

The Software Scan - Ignore Software table is displayed.



- 2 You can add entries to the Ignore Software table either manually or automatically.

Manually Adding Entries to the Table

1. Click *Insert* to add a new row.

2. In the *Pattern Type* drop-down list, select *Expandable Expression* or *Regular Expression*.
3. Specify a software name.
4. Click *OK*.

For example, if you do not want the Inventory scanner to scan for the Adobe products, configure the following settings:

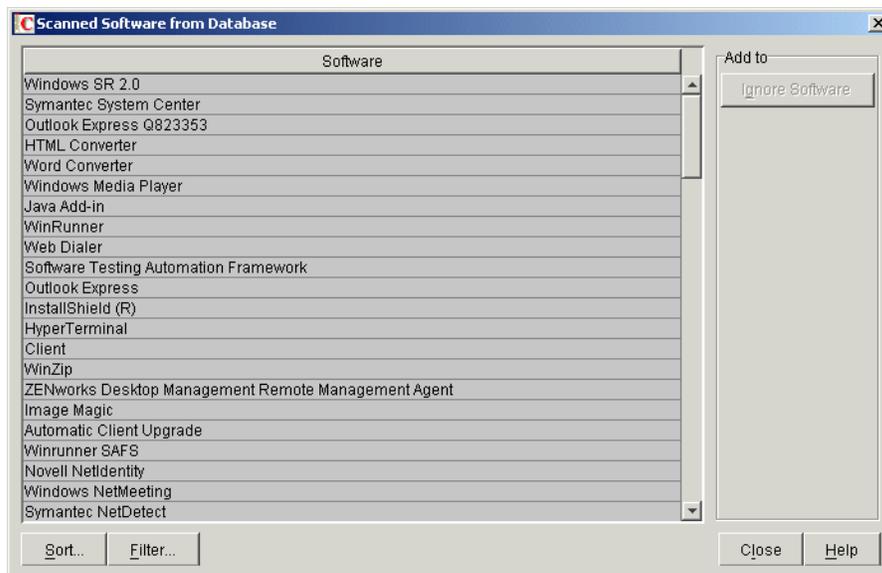
Pattern Type= Expandable Expression

Software = Adobe*

The Inventory scanner does not report the software that has names beginning with Adobe.

Automatically Adding Entries to the Table

1. Click *Database* located in the *Add From* pane.
The Scanned Software from Database dialog box is displayed.



2. Select the software that you want to add to the Ignore Software table.
3. Click the *Ignore Software* button located in the *Add to* pane.
4. Click *Close*.

IMPORTANT: The table might contain **inherited rules**. You can edit or delete these rules only in the software dictionary of the inventory server configured in the Dictionary Update policy. These rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

You can also perform the following operations in the Software Scan - Ignore Software table:

- ◆ Deleting only the non-inherited entries.
- ◆ **Sorting Entries in the Table.**
- ◆ **Filtering Entries in the Table.**
- ◆ **Refreshing Entries in the Table.**

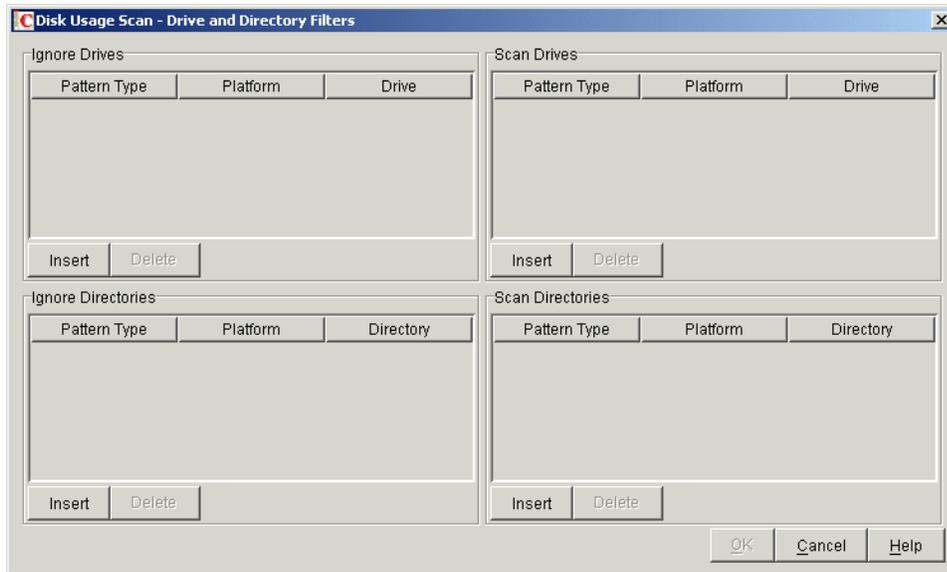
16.3.21 Disk Usage Scanning Filters - Drives and Directories

The “Drives and Directories” filter allows you to configure drives and directories to be included or excluded during the Inventory disk usage scanning.

To configure this filter:

- 1 In the *Filters* property page, click the *Edit Table* option of *Drives and Directories* located in the *Disk Usage Scanning Filters* pane.

The Disk Usage Scan - Drive and Directory Filters dialog box is displayed.



- 2 Configure the following filters:

- ♦ “Ignore Drives” on page 619
- ♦ “Scan Drives” on page 619
- ♦ “Ignore Directories” on page 620
- ♦ “Scan Directories” on page 621

By default, the Inventory scanner scans the disk usage of all directories on the inventoried servers. If you have configured all directories to be ignored during a disk usage scan using the Ignore Directories filter, but now want to include a specific directory in scan, identify the specific directory in the Scan Directories filter. The settings of the Scan Directories filter override the settings of the Ignore Directories and Ignore Drives filters.

For example, if you want the Inventory scanner to ignore the disk usage of all files and directories in C: except for the `c:\program files` directory on Windows inventoried servers, configure the following filters as shown below:

- ♦ **Ignore Drives:** Configure the following settings:
 - Pattern Type = System Expandable Expression
 - Platform = Windows
 - Drive = C
- ♦ **Scan Directories:** Configure the following settings:

Pattern Type = System Expandable Expression
Platform= Windows
Drive=c:\program files

3 Click *OK*.

Ignore Drives

The “Ignore Drives” filter allows you to specify the drives that should not be scanned for disk usage at the inventoried servers.

By default, the Inventory scanner scans all drives.

To configure the “Ignore Drives” filter:

- 1 In the Ignore Drives table, click *Insert* to add a new row.
- 2 In the *Pattern Type* drop-down list, select *System Expandable Expression*, *Expandable Expression*, or *Regular Expression*.
- 3 (Conditional) If you select *System Expandable Expression* as the pattern type, then select *NetWare* or *Windows* in the *Platform* drop-down list, depending on the operating system of the inventoried servers.

IMPORTANT: If you select *Expandable Expression* or *Regular Expression* as the pattern type, the corresponding value in the *Platform* column automatically changes to *Any*. You cannot change the value.

- 4 Specify a drive name.

For example, if you want the Inventory scanner not to scan for the disk usage of C drive on all the Windows* inventoried servers, configure the following settings in the Ignore Drives table:

Pattern Type = System Expandable Expression
Platform = Windows
Drive = C

The Inventory scanner does not scan the disk usage of files on the C drive.

IMPORTANT: The table might contain **inherited rules**. You can edit or delete these rules only in the software dictionary of the inventory server configured in the Dictionary Update policy. These rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

To delete an entry from the Ignore Drives table, select the entry and click *Delete*. You can delete only the non-inherited entries.

Scan Drives

The “Scan Drives” filter allows you to specify the drives whose disk usage should be scanned for at the inventoried servers.

To configure the “Scan Drives” filter:

- 1 In the Scan Drives table, click *Insert* to add a new row.

- 2 In the *Pattern Type* drop-down list, select *System Expandable Expression*, *Expandable Expression*, or *Regular Expression*.
- 3 (Conditional) If you select *System Expandable Expression* as the pattern type, then select *NetWare* or *Windows* in the *Platform* drop-down list, depending on the operating system of the inventoried servers.

IMPORTANT: If you select *Expandable Expression* or *Regular Expression* as the pattern type, the corresponding value in the *Platform* column automatically changes to *Any*. You cannot change the value.

- 4 Specify a drive name.

For example, if you want the Inventory scanner to scan for the disk usage of C drive on all the Windows inventoried servers, configure the following settings in the Scan Drives table:

Pattern Type = System Expandable Expression
Platform = Windows
Drive = C

You must also configure the following settings in the Ignore Drives table:

Pattern Type = System Expandable Expression
Platform = Windows
Drive = *

The Inventory scanner scans and stores the disk usage of the files in the C drive into the Inventory database.

IMPORTANT: The table might contain **inherited rules**. You can edit or delete these rules only in the software dictionary of the inventory server configured in the Dictionary Update policy. These rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

To delete an entry from the Scan Drives table, select the entry and click *Delete*. You can delete only the non-inherited entries.

IMPORTANT: By default, the Inventory scanner scans the disk usage of all drives on the inventoried servers. If you have configured all drives to be ignored during a disk usage scan using the Ignore Drives filter, but now want to include a specific drive in the scan, identify the specific drive in the Scan Drives filter. The settings of the Scan Drives filter override the settings of the Ignore Drives filter.

Ignore Directories

The “Ignore Directories” filter allows you to specify the directories whose disk usage should not be scanned for at the inventoried servers.

By default, the Inventory scanner scans all directories.

To configure the “Ignore Directories” filter:

- 1 In the Ignore Directories table, click *Insert* to add a new row.

- 2 In the *Pattern Type* drop-down list, select *System Expandable Expression*, *Expandable Expression*, or *Regular Expression*.
- 3 (Conditional) If you select *System Expandable Expression* as the pattern type, then select *NetWare* or *Windows* in the *Platform* drop-down list, depending on the operating system of the inventoried servers.

IMPORTANT: If you select *Expandable Expression* or *Regular Expression* as the pattern type, the corresponding value in the *Platform* column automatically changes to *Any*. You cannot change the value.

- 4 Specify a directory name.

For example, if you want the Inventory scanner not to scan for the disk usage of the `c:\program files` directory on all the Windows inventoried servers, configure the following settings in the Ignore Directories table:

Pattern Type = System Expandable Expression

Platform = Windows

Directory= `c:\program files`

The Inventory scanner does not scan for the disk usage of `c:\program files`.

IMPORTANT: The table might contain **inherited rules**. You can edit or delete these rules only in the software dictionary of the inventory server configured in the Dictionary Update policy. These rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

To delete an entry from the Ignore Directories table, select the entry and click *Delete*. You can delete only the non-inherited entries.

Scan Directories

The “Scan Directories” filter allows you to specify the directories whose disk usage should be scanned for at the inventoried servers.

To configure the “Scan Directories” filter:

- 1 In the Scan Directories table, click *Insert* to add a new row.
- 2 In the *Pattern Type* drop-down list, select *System Expandable Expression*, *Expandable Expression*, or *Regular Expression*.
- 3 (Conditional) If you select *System Expandable Expression* as the pattern type, then select *NetWare* or *Windows* in the *Platform* drop-down list, depending on the operating system of the inventoried servers.

IMPORTANT: If you select *Expandable Expression* or *Regular Expression* as the pattern type, the corresponding value in the *Platform* column automatically changes to *Any*. You cannot change the value.

- 4 Specify a directory name.

For example, if you want the Inventory scanner to scan for disk usage of the `c:\program files` directory on all the Windows inventoried servers, configure the following settings in the Scan Directories table:

Pattern Type = System Expandable Expression
Platform = Windows
Directory = c:\program files

You must also configure the following settings in the Ignore Directories table:

Pattern Type = System Expandable Expression
Platform = Windows
Directory = *

The Inventory scanner scans and stores only disk usage of files in c:\program files into the Inventory database.

IMPORTANT: The table might contain **inherited rules**. You can edit or delete these rules only in the software dictionary of the inventory server configured in the Dictionary Update policy. These rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

To delete an entry from the Scan Directories table, select the entry and click *Delete*. You can delete only the non-inherited entries.

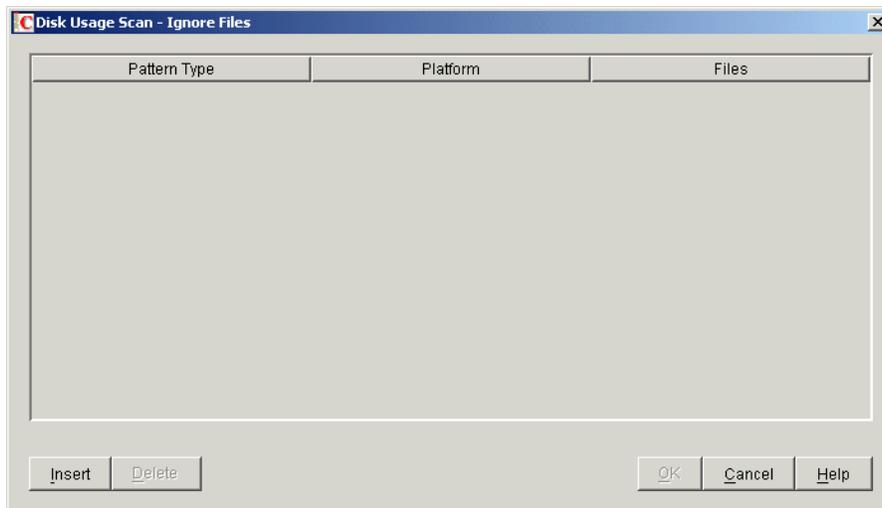
16.3.22 Disk Usage Scanning Filters - Files

The “Files” filter allows you to configure files to be excluded during the Inventory disk usage scanning.

To configure this filter:

- 1 In the Filters property page, click the *Edit Table* option of *Files* located in the *Disk Usage Scanning Filters* pane.

The Disk Usage Scan - Ignore Files dialog box is displayed.



- 2 Click *Insert* to add a new row.

- 3 In the *Pattern Type* drop-down list, select *System Expandable Expression*, *Expandable Expression*, or *Regular Expression*.
- 4 (Conditional) If you select *System Expandable Expression* as the pattern type, then select *NetWare* or *Windows* in the *Platform* drop-down list, depending on the operating system of the inventoried serves.

IMPORTANT: If you select *Expandable Expression* or *Regular Expression* as the pattern type, the corresponding value in the *Platform* column automatically changes to *Any*. You cannot change the value.

- 5 Specify a file.
- 6 Click *OK*.

For example, if you want the Inventory scanner to scan for disk usage of all files with extension `.exe`, except `msoffice.exe`, configure the following rules as shown below:

- ♦ **Disk Usage Scan - Ignore Files:** Configure the following settings:
Pattern Type = Expandable Expression
Files = `msoffice.exe`
- ♦ **Report Disk Space used by file extensions:** Configure the following settings:
Pattern Type = Expandable Expression
Files = `exe`

IMPORTANT: The table might contain **inherited rules**. You can edit or delete these rules only in the software dictionary of the inventory server configured in the Dictionary Update policy. These rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

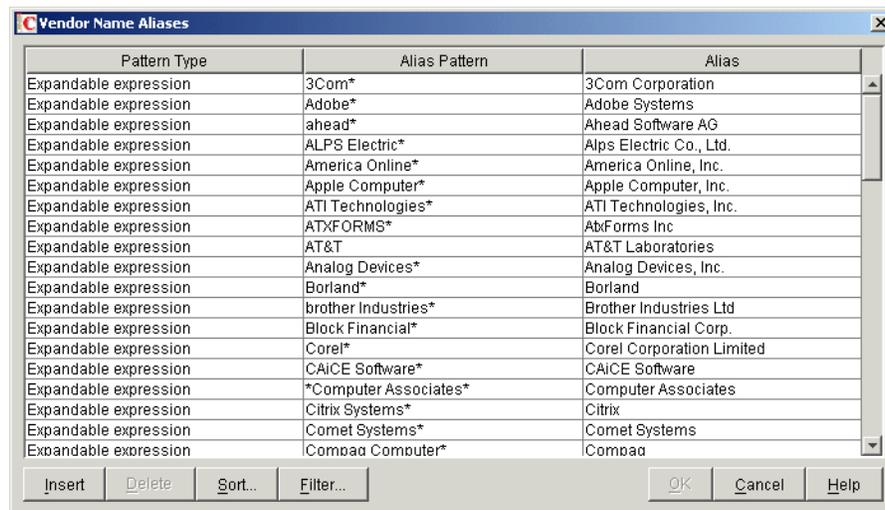
To delete an entry from the table, select the entry and click *Delete*. You can delete only the non-inherited entries.

16.3.23 Vendor Name Aliases

The “Vendor Name Aliases” rule allows you to configure aliases for vendor names.

- 1 In the *Aliases* property page, click the *Edit Table* option of *Vendor Name Aliases* located in the *Specify Aliases* pane.

The Vendor Name Aliases table is displayed.



| Pattern Type | Alias Pattern | Alias |
|-----------------------|-----------------------|---------------------------|
| Expandable expression | 3Com* | 3Com Corporation |
| Expandable expression | Adobe* | Adobe Systems |
| Expandable expression | ahead* | Ahead Software AG |
| Expandable expression | ALPS Electric* | Alps Electric Co., Ltd. |
| Expandable expression | America Online* | America Online, Inc. |
| Expandable expression | Apple Computer* | Apple Computer, Inc. |
| Expandable expression | ATI Technologies* | ATI Technologies, Inc. |
| Expandable expression | ATXFORMS* | AbForms Inc |
| Expandable expression | AT&T | AT&T Laboratories |
| Expandable expression | Analog Devices* | Analog Devices, Inc. |
| Expandable expression | Borland* | Borland |
| Expandable expression | brother Industries* | Brother Industries Ltd |
| Expandable expression | Block Financial* | Block Financial Corp. |
| Expandable expression | Corel* | Corel Corporation Limited |
| Expandable expression | CAICE Software* | CAICE Software |
| Expandable expression | *Computer Associates* | Computer Associates |
| Expandable expression | Citrix Systems* | Citrix |
| Expandable expression | Comet Systems* | Comet Systems |
| Expandable expression | Compaq Computer* | Compaq |

- 2 Click *Insert* to add a new row.
- 3 In the *Pattern Type* drop-down list, select *Expandable Expression* or *Regular Expression*.
- 4 Specify an alias pattern.
- 5 Specify an alias.
- 6 Click *OK*.

For example, if you want the Inventory scanner to report all instances of the vendor name beginning with “Microsoft” as “Microsoft Corporation” in the Inventory database, configure the following settings:

Pattern Type = Expandable Expression

Alias Pattern = Microsoft*

Alias = Microsoft Corporation

If the Inventory scanner reports Microsoft, Microsoft Inc., or Microsoft Inc. Corporation vendor names during the scan, then the name of the vendor beginning with “Microsoft” is stored as “Microsoft Corporation” in the Inventory database.

IMPORTANT: The table might contain **inherited rules**. You can edit or delete these rules only in the software dictionary of the inventory server configured in the Dictionary Update policy. These rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

You can also perform the following operations in the Vendor Name Aliases table:

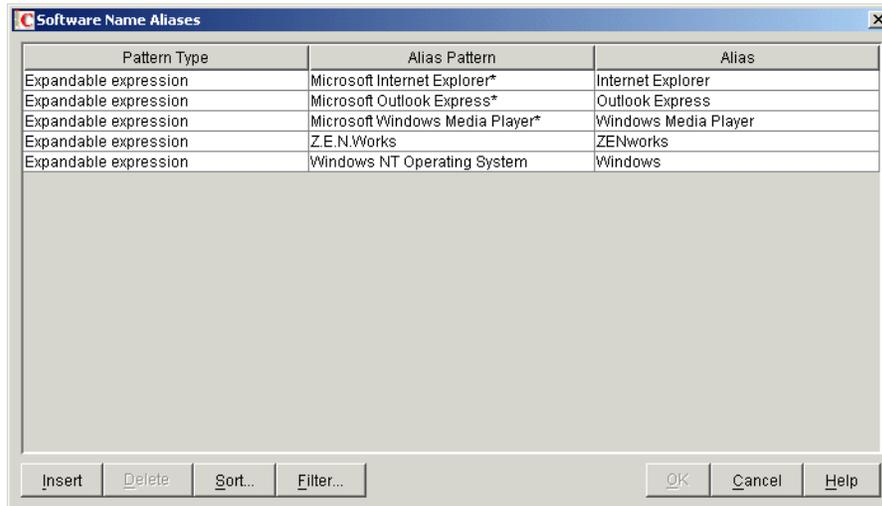
- ◆ Deleting only the non-inherited entries.
- ◆ **Sorting Entries in the Table.**
- ◆ **Filtering Entries in the Table.**
- ◆ **Refreshing Entries in the Table.**

16.3.24 Software Name Aliases

The “Software Name Aliases” rule allows you to configure aliases for software names.

- 1 In the *Aliases* property page, click the *Edit Table* option of *Software Name Aliases* located in the *Specify Aliases* pane.

The Software Name Aliases table is displayed.



| Pattern Type | Alias Pattern | Alias |
|-----------------------|---------------------------------|----------------------|
| Expandable expression | Microsoft Internet Explorer* | Internet Explorer |
| Expandable expression | Microsoft Outlook Express* | Outlook Express |
| Expandable expression | Microsoft Windows Media Player* | Windows Media Player |
| Expandable expression | Z.E.N.Works | ZENworks |
| Expandable expression | Windows NT Operating System | Windows |

- 2 Click *Insert* to add a new row.
- 3 In the *Pattern Type* drop-down list, select *Expandable Expression* or *Regular Expression*.
- 4 Specify an alias pattern.
- 5 Specify an alias.
- 6 Click OK.

For example, if you want the Inventory scanner to report all instances of the product name “WinZip” as “WinZip Application” in the Inventory database, configure the following settings:

Pattern Type = Expandable Expression

Alias Pattern = WinZip

Alias = WinZip Application

If the Inventory scanner scans the WinZip, WinZip Executables, or WinZip Applications product names, then the name of the software that exactly matches “WinZip” is stored as “WinZip Application” in the Inventory database. The remaining software names are reported as scanned.

IMPORTANT: The table might contain **inherited rules**. You can edit or delete these rules only in the software dictionary of the inventory server configured in the Dictionary Update policy. These rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

You can also perform the following operations in the Software Name Aliases table:

- ♦ Deleting only the non-inherited entries.
- ♦ **Sorting Entries in the Table.**

- ◆ **Filtering Entries in the Table.**
- ◆ **Refreshing Entries in the Table.**

16.3.25 Reconcile Software

The “Reconcile Software” rule allows you to associate the software identified through Add/Remove Programs or the MSI, with an appropriate software and vendor identified and configured through the ZENworks software dictionary. The association might be necessary because the software entries in Add/Remove Programs or the MSI might not use the same software name and vendor as configured in the ZENworks software dictionary.

To configure the rule:

- 1 In the *Aliases* property page, click the *Edit Table* option of *Reconcile Software* located in the *Reconcile Software* pane.

The Reconcile Software table is displayed.

| Add-Remove Program(ARP/MSI Software name) | Displayed ARP/MSI Software name | Software name | Vendor |
|---|---------------------------------|-----------------------|---------------------------|
| {624C9AE0-6CD8-4166-9D... | | XML Spy | Altova |
| Adobe Acrobat 5.0 | | Acrobat Reader | Adobe Systems |
| Adobe Acrobat Reader 3.02 | | Acrobat Reader | Adobe Systems |
| LiveUpdate1.6 | | LiveUpdate | Symantec Corporation |
| LiveUpdate1.7 | | LiveUpdate | Symantec Corporation |
| Visual C++ 6.0 Professional... | | Visual C++ | Microsoft |
| {A4D7B764-4140-11D4-88... | | Nero - Burning Rom | Ahead |
| TextPad 4 | | TextPad | Helios Software Solutions |
| Winamp | | Winamp | America Online |
| Winamp3 | | Winamp | America Online |
| AR System User 5.1 | | Action Request System | BMC Software company |
| AR System User 5.1 | | Action Request System | BMC Software company |
| CM Synergy 6.2 | | CM Synergy | Telelogic AB |
| CONSOLE1 | | ConsoleOne | Novell |
| {7699B723-9718-41DE-8C... | | Crystal Reports | Seagate |
| Java 2 SDK Standard Editio... | | Java2 SDK | Sun |
| Java 2 SDK Standard Editio... | | Java2 SDK | Sun |
| {35A3A4F4-B792-11D6-A78... | | Java2 SDK | Sun |
| JRE 1.3.1_01 | | Java JRE | Sun |

By default, the Reconcile Software table displays pre-defined mapping of software in the Add/Remove Programs or MSI with the software configured in the software dictionary. It also displays the Add/Remove Programs or MSI software identified during the last scan for which you can configure software and vendor names.

This table has the following columns:

- ◆ The Add-Remove Program (ARP) key.
You cannot edit the values of this attribute.
- ◆ The ARP /MSI name as displayed either in Add/Remove Programs or in the MSI.
The Displayed ARP/MSI Software name shows the software identified through Add/Remove Programs or the MSI and stored in the Inventory database.
You cannot edit the values of this attribute.
- ◆ The software name associated with its corresponding Add/Remove Programs or MSI name.
- ◆ The vendor name associated with its corresponding Add/Remove Programs or MSI name.

IMPORTANT: The table might contain **inherited rules**. You can edit or delete these rules only in the software dictionary of the inventory server configured in the Dictionary Update policy. These rules are displayed in dark gray color if you are accessing the software dictionary of ZENworks 7 Inventory server.

- 2 To reconcile software, do the following in this table:
 - ♦ Specify software and vendor names for the software identified through Add/Remove Programs or the MSI but not yet been associated.
 - ♦ (Optional) Change the software and vendor names for the software that has already been configured in the software dictionary.

You can also perform the following operations in the Reconcile Software table:

- ♦ Deleting only the non-inherited entries.
- ♦ **Sorting Entries in the Table.**
- ♦ **Filtering Entries in the Table.**
- ♦ **Refreshing Entries in the Table.**

16.3.26 Sorting Entries in the Table

You can sort the entries in the table by one, two, or three columns.

- 1 Click *Sort*.

The Sort dialog box is displayed.

- 2 In the *Sort by* drop-down list, select the table column by which you want to sort the entries.
- 3 Select *Ascending* or *Descending*.
- 4 (Optional) To sort by either two or three columns, configure the *Then by* drop-down lists and select *Ascending* or *Descending*.
- 5 Click *OK*.

16.3.27 Filtering Entries in the Table

- 1 Click *Filter*.

The Filter dialog box is displayed.

- 2 In the Filter dialog box, do the following to create a query:
 - 2a Select an attribute.
 - 2b Select an operator. The operators displayed depend on the attribute you select in Step 2a.
 - 2c Type a value.
 - 2d (Optional) To create an advanced query, select one of the following logical operators and define the query:

| Logical Operator | Functionality |
|------------------|---|
| AND | Creates a new row. The filter displays items that match the conditions in each row joined by AND. |

| Logical Operator | Functionality |
|------------------|---|
| OR | Creates a new row. The filter displays items that match the conditions in either row joined by OR. |
| New Row | Creates a new row to form a new query. |
| Delete Row | Deletes the row from the filter. |
| End | Closes the query expression. If you select End in a row that is followed by other rows, the subsequent rows and groups are deleted. |

3 Click *OK*.

After applying the filter, the table list displays only the resulting entries. To clear the filter:

1 Click *Filter*.

The Filter dialog box is displayed.

2 Click *Clear*, then click *OK*.

16.3.28 Refreshing Entries in the Table

Use the *Refresh* option if you want to reapply the sort or filter operations. To refresh entries in the table, either click *Refresh* or press F5.

IMPORTANT: The *Refresh* button is displayed only when you apply sort or filter operations to the table.

16.3.29 Disabling File Scan

You can disable the software scanning for all software except for the following:

- ◆ Software registered in the Add/Remove Programs dialog box
- ◆ Software installed through MSI
- ◆ Software scanned by default such as Microsoft Windows, Internet Explorer, Outlook, MediaPlayer, ZENworks, Novell client, Microsoft Office and the set of Antivirus programs

To disable the file scanning:

1 In ConsoleOne, right-click the Inventory Service Object, click *Properties*.

2 Click the *Software Inventory Configuration* tab.

The Software Configuration page is displayed by default.

3 Select the *Ignore Default File-Software Mapping Rules* check box.

4 Click the *Edit Table* option of the “Report Files with These File Extensions As Unidentified Software” rule.

5 By default, the table has an entry with the EXE file extension. Delete the entry.

6 Click *OK*.

7 Click *Apply*.

8 Click *Edit Table* option of the “Report Disk Space used by File Extensions”.

- 9 By default, the table has entries with EXE, DLL, MP3, and JPEG file extensions. Delete all the entries.
- 10 Click *OK*.
- 11 Click *Apply*, then click *Close*.

16.3.30 Base-lining the Software Dictionary Deployment

When you deploy the software dictionary for the first time, the default dictionary settings are effective and the Inventory scanner reports the following information:

- ◆ Unidentified software
- ◆ Multiple instances of software installed on the same inventoried server
- ◆ Redundant drives and directories that do not contain software

This scan results in storing huge amount of irrelevant information in the Inventory database. It also degrades the performance of Storer and all Inventory ConsoleOne utilities such as Query, Reporting, etc.

To avoid these problems, we recommend you to fine tune the software dictionary according to your requirements before deploying it in your enterprise. Do the following to fine tune the software dictionary:

- 1 Deploy a small representative set of inventoried servers in the test lab.

NOTE: This representative set should typically represent all sections or departments of your enterprise that you want to collect inventory for.

- 2 Attach these inventoried servers to a Standalone Inventory Server, which is also deployed in the test lab.
- 3 In addition to defaults that are listed in the table of the **Report Files with These File Extensions As Unidentified Software** rule, you may want to scan for additional file extensions and report them as identified software. These could be extensions of application files like DLLs, etc.
- 4 Schedule the scan and wait until the inventory information is stored into the Inventory database.
- 5 Re-configure the software dictionary based on the inventory information that is available in the database to resolve the above discussed problems. Perform the following tasks:
 - ◆ **UnIdentified Software:** Based on the **Report Files with These File Extensions As Unidentified Software** settings, all the information related to the unidentified software can be viewed in the **Manage Unidentified Software** table.

The result will contain the following:

- ◆ Applications that are not yet identified by the software dictionary.
- ◆ Application files that are already identified by the software dictionary.
- ◆ Application files that might be redundant such as Operating System files or DOS files.

Perform the following tasks in the Manage Unidentified Software table:

- ◆ Add the applications that are not yet identified by the software dictionary to the **Software Dictionary** table using the *Software Dictionary* button located in the *Add To* pane.

- ♦ Add the application files that are part of already identified by the software dictionary and application files that might be redundant to the [Software Scanning Filters - Files](#) table using the Ignore Files button located in the Add To pane.

The effectiveness of this exercise is based on the following assumptions:

- ♦ The representative set should not be different from the sections or departments of your enterprise; otherwise it would amount for large number of un-identified software being scanned and reported.
 - ♦ The inventoried servers in the enterprise are largely controlled by the enterprise administrator, who installs and copies the non-standard applications.
 - ♦ **Multiple instances of Software on the same inventoried server:** For an inventoried server, the same software can be reported twice if one entry is reported from the Add Remove Program scanning or the MSI scanning, and the other is reported based on the software dictionary configuration. The [Section 16.3.25, “Reconcile Software,” on page 626](#) rule contains default configurations to merge these two entries but this may not be complete. In order to resolve this problem, you must manually configure the Edit Add-Remove Software rule.
 - ♦ **Redundant drives and directories that do not contain software:** Configure the rules in Software Scanning page and the Disk Usage Scanning pages of the software dictionary to eliminate these drives and directories from scan. For more information about the software dictionary rules, see [Step 3 on page 597](#).
- 6 Re-scan all the inventoried servers.
 - 7 After the inventory information is stored in the Inventory database, you could notice that all the entries that you marked for dictionary during the earlier scan would be scanned and reported as a software.
 - 8 Repeat Step 3 through Step 7 till you fine tune the dictionary according to your requirements.

16.3.31 Viewing Software Information in the Inventory Summary

- 1 In ConsoleOne, configure the Inventory database. For more information on how to configure the database, see [“Configuring the Inventory Database” on page 635](#)
- 2 Right-click an inventoried server, click *Actions*, then click *Inventory*.
- 3 In the Summary dialog box, click *Inventory Information > Hardware/Software Inventory > Software > Application Vendors* to view the software inventory information.

A list of Software Group and Software of the vendor is displayed. Software Group includes software patch and representative file information of the group. Software includes software patch and representative file information of the product.

For more information, see [“Viewing the Inventory Summary of an Inventoried Server” on page 636](#)

16.3.32 Generating Software Inventory Reports

You can now generate the following Software Inventory reports:

- ♦ Add-Remove Programs by Application
- ♦ Add-Remove Programs by Machine

- ◆ Anti-Virus Signature Files by Machine
- ◆ Anti-Virus Signature Machine Count
- ◆ Disk Usage by Machine
- ◆ Exception List by Machine
- ◆ Installed NetWare Software by Machine
- ◆ Internet Explorer Installation Count
- ◆ Internet Explorer Patches by Machine
- ◆ Internet Explorer by Machine
- ◆ MSI Products by Application
- ◆ MSI Products by Machine
- ◆ Microsoft Office Components by Machine
- ◆ Microsoft Office Installation Count
- ◆ Microsoft Office by Machine
- ◆ Novell Client Components by Machine
- ◆ Novell Client Installation Count
- ◆ Novell Client by Machine
- ◆ Novell ZENworks Desktop Management Installed Agent Components by Machine
- ◆ Novell ZENworks Desktop Management Installed Server Components by Machine
- ◆ Novell ZENworks Handheld Management Installed Components by Machine
- ◆ Novell ZENworks Installed Components by Machine
- ◆ Novell ZENworks Installed Suites by Machine
- ◆ Novell ZENworks Server Management Installed Agent Components by Machine
- ◆ Novell ZENworks Server Management Installed Server Components by Machine
- ◆ Outlook Express Installation Count
- ◆ Outlook Express by Machine
- ◆ Software Dictionary Application Files by Machine
- ◆ Software Dictionary Applications by Machine
- ◆ Software Dictionary Versions Machine Count
- ◆ Software Dictionary Versions by Machine
- ◆ Software Installation Count
- ◆ Software Installations
- ◆ Software by Machine
- ◆ System Software Inventory Report
- ◆ Windows Components by Machine
- ◆ Windows Installation Count
- ◆ Windows Media Player Count
- ◆ Windows Media Player Patches by Machine
- ◆ Windows Media Player by Machine

- Windows Operating System by Machine
- Windows Security Patches by Machine
- Windows Security Patches by Patch

For more information about each report, see [“Types of Inventory Reports” on page 652](#).

16.4 Customizing the Software Inventory Information To Be Scanned For ZENworks for Servers 3.x Inventoried Servers

Refer to the [ZENworks for Servers 3.0.2 Documentation Web site \(http://www.novell.com/documentation/zfs302/index.html\)](http://www.novell.com/documentation/zfs302/index.html) to know how to customize the software inventory information for the ZENworks for Servers 3.x inventoried servers.

16.5 Removing Redundant Inventoried Servers from the Inventory Database

You can remove the unwanted, redundant, or obsolete inventoried servers from the Inventory database using the Inventory Removal service.

The Inventory Removal service is a manual service that runs on the Inventory server. The service removes the inventoried servers from the Inventory database using the `inventoryremovallist.txt` file, which contains a list of inventoried servers that must be removed from the Inventory database.

IMPORTANT: You can run the Inventory Removal service on the Intermediate Server only if the Intermediate Server has either inventoried servers or database attached to it.

To remove the inventoried servers from the Inventory database:

- 1 Using a text editor, create a file with the name `inventoryremovallist.txt` with the following contents:

```

;                               Enter comments, if any
DN of the inventoried server (as stored in the Inventory database)
to be removed from the Inventory database
DN of the inventoried server (as stored in the Inventory database)
to be removed from the Inventory database
....
DN of the inventoried server (as stored in the Inventory database)
to be removed from the Inventory database

```

A sample `inventoryremovallist.txt` file is as follows:

```

CN=INT-SERVER-NDS.OU=Leaf.O=XYZ.T=XYZ-TREEzen-server.xyz.com
CN=ROOT-SERVER-NDS.O=XYZ.T=XYZ-TREE

```

To generate the list of inventoried servers that must be removed, you can either perform a query on a selected criteria or manually enter the names of the inventoried servers. For more information on Query, see [“Viewing Inventory Information of Inventoried Servers by Querying the Database” on page 649](#).

- 2** Copy the `inventoryremovallist.txt` file to the `ZENworks_installation_path\zenworks\inv\server\wminv\properties` directory.
- 3** In the `ZENworks_installation_path\zenworks\inv\server\wminv\properties\inventoryremoval.properties` file, ensure that the value of `FilePath` is the location of `inventoryremovallist.txt` (specified in **Step 2**).

NOTE: Ensure that the path separator is a forward slash (/) and not a backslash (\).

- 4** At the server console prompt, enter `StartSer RemoveInventory` to start the Inventory Removal service.

The Inventory Removal service operates in the following order:

- 1** The Inventory Removal service reads each line of the `inventoryremovallist.txt` file and creates a `delete str` file for each inventoried server that is listed in the `inventoryremovallist.txt` file.

The `delete str` file is saved in the `scandir` directory if the Selector is running, else it will be placed in the `dbdir` or `entmergedir` directories depending on the Inventory server role.

- 2** The Selector validates the `delete str` file and copies it into the `dbdir` and `entmergedir` directories.
- 3** The Storer reads the `delete str` file from `dbdir` and deletes the inventoried server from the attached Inventory database.
- 4** If the inventory deployment rolls up inventory information, the `delete str` is also rolled up to the next level Inventory server.

The inventoried server is deleted from the Inventory database at all Inventory servers deployed at the enterprise level.

The following sections indicate how you can view the inventory information:

- ◆ [Section 17.1, “Viewing the Inventory Information Using ConsoleOne,” on page 635](#)
- ◆ [Section 17.2, “Exporting the Inventory Information,” on page 676](#)
- ◆ [Section 17.3, “Retrieving Inventory information from the Inventory Database Without Using the CIM Schema,” on page 683](#)

17.1 Viewing the Inventory Information Using ConsoleOne

The following sections explain the various types of information you can view using ConsoleOne:

- ◆ [Section 17.1.1, “Configuring the Inventory Database,” on page 635](#)
- ◆ You can list hardware and software components found on the inventoried server and any custom information you have specified for the inventoried server.
The Inventory Summary window displays the inventory items for an inventoried server. This window displays the information from the last inventory scan for the inventoried server. For more information, see [Section 17.1.2, “Viewing the Inventory Summary of an Inventoried Server,” on page 636](#).
- ◆ You can list inventoried servers with the inventory information from the Inventory database satisfying the criteria you specify in the Inventory Query window. You form a query by specifying the component and its attribute for servers within the selected database sites.
For more information about querying the Inventory database, see [Section 17.1.3, “Viewing Inventory Information of Inventoried Servers by Querying the Database,” on page 649](#).
- ◆ You can use a list of reports that generate the inventory information from the Inventory database specific to your needs.
For more information, see [Section 17.1.4, “Running Inventory Reports,” on page 651](#).
- ◆ You can now quickly and easily view the inventory information
For more information, see [Section 17.1.5, “Quickly and Easily Viewing the Inventory Data Using Quick Reports,” on page 664](#).

17.1.1 Configuring the Inventory Database

If you want to view the inventory information stored in the database from ConsoleOne, you must configure the database. The inventory information from the Inventory database that you configure is used for generating inventory reports, viewing inventory information, and for querying the inventory information from the database.

To configure the Inventory database:

- 1 In ConsoleOne, select a container.

- 2 Invoke Configure DB.
 - ♦ To invoke Configure DB from a database object, right-click the database object, click *ZENworks Inventory*, then click *Configure DB*. This configures the database object.
 - ♦ To invoke the Configure DB dialog box from the ConsoleOne Tools menu, click *Tools*, click *ZENworks Inventory*, then click *Configure DB*.
- 3 Click *Browse* to browse for and select the *ZENworks Database* object.

You can also select an existing ZENworks Database object from the list of Database objects. This Database object contains the database settings such as the protocol, port in use by the database, and others.
- 4 To apply this database configuration to all the sessions, select the *Apply Configuration Across Sessions* check box.
- 5 Click *OK*.

The database you configured is used for data retrieval unless you change it again using this same procedure.

17.1.2 Viewing the Inventory Summary of an Inventoried Server

The Inventory Summary window displays the information from the last inventory scan for the inventoried server.

To view the inventory information of an inventoried server, do the following in ConsoleOne

- 1 Configure the Inventory database.

For more information, see [“Configuring the Inventory Database” on page 635](#).
- 2 Right-click any of the following objects: Subscriber, Distributor, or External Subscriber, click *Actions*, then click *Inventory*.

or

In the Query Results window, double-click an inventoried server.

ZENworks 7 Server Management provides the following inventory information collected from the inventoried servers:

Table 17-1 *Inventory information as displayed in Inventory Summary*

| Scan Data Group | Scan Data Item | Description |
|---|------------------|--|
|  Inventory Information | Inventory Server | Name of the Inventory server to which the scans are sent |
| | Last Scan Date | List of all inventoried servers that were scanned on or before the specified date and time |
| | Scan Mode | Mode used by the Inventory scanner to scan the inventoried server |
| | Version | Version number of the Inventory scanner |

| Scan Data Group | Scan Data Item | Description |
|--|----------------------------|--|
| | General Dictionary Version | Version number of the General Dictionary NOTE: The General Dictionary version is not same as the ZENworks product version. |
| | Private Dictionary Version | Version number of the Private Dictionary |
|  Hardware/Software Inventory >  General >  System Information | Asset Tag | Asset tag number that the ROM-based setup program creates |
| | Computer Model | Identifying information of the computer such as Compaq or Dell |
| | Computer Type | Type of computer such as IBM PC |
| | Machine Name | DNS name of the inventoried server |
| | Management Technology | Technology available on the inventoried server such as DMI, WMI, and others |
| | Model Number | Model number of the computer |
| | Serial Number | Serial number of the computer system assigned by manufacturer |
| | Tag | Unique identifier of system information |
|  Hardware/Software Inventory >  General >  System Identification | Primary Owner Name | The name of the primary user or owner of this system |
| | Primary Owner Contact | The phone number of the primary user of this system |
| | Name | Name of the inventoried server as represented in eDirectory, such as the fully qualified DN of the inventoried server |
|  Hardware/Software Inventory >  General >  Login Details >  eDirectory Login Details | Current login user | User logged in to the Primary eDirectory tree when the inventoried server was scanned |
| | Last login user | User most recently logged in to the Primary eDirectory tree through Novell Client when the inventoried server was scanned |
|  Hardware/Software Inventory >  General >  Login Details >  Windows Domain | Name | Domain name of the inventoried server |
|  Hardware/Software Inventory >  Software >  Application Vendors > <i>Vendor_name</i> >  <i>software_group_name</i> >  <i>software</i> | Name | Vendor-defined name of the product represented as a vendor trademark or registered trademark |

| Scan Data Group | Scan Data Item | Description |
|-----------------|-------------------------|--|
| | Version | User-friendly version of a product. For example, the version for Windows 2000 is 2000 or Major.Minor Version of the Product |
| | Category | Product category to which the product belongs For example, Office is a part of the Productivity tools category and Solitaire is a game |
| | Description | Description of the product |
| | Help Link | Support Web site URL for the product that is available in ARP and MSI |
| | Package GUID | Vendor-defined GUID for a product that is available in MSI |
| | Product Identifier | A unique, 16-character identifier for an installed product. This identifier is available from MSI on Windows. The format is ABCD-1234-WXYZ-PQRS |
| | Internal Version | Internal version of a product The format is: <i>major version.minor version.build.sub build number</i> or <i>major version.minor version.build</i> |
| | Language | User-friendly name for the language of this copy of the product |
| | Uninstall String | The command to invoke for uninstalling this product instance. Currently, this is available in Add/Remove Programs (ARP) and MSI on Windows. |
| | Install Source | Identifies the file system path where the installation files were stored when installing this product instance. Currently, this is available in ARP and MSI on Windows. |
| | Last Execution Time | Date and time stamp when the product was last executed |
| | Frequency of Usage | Number of times the product has been used |
| | Friendly Name | Display name of the software |
| | Installation Repository | Source of scan, which can be Add/Remove Programs, MSI, Software Dictionary, or PRODUCTS.DAT |

| Scan Data Group | Scan Data Item | Description |
|---|------------------------------|--|
| | Support Pack | Installed support pack number of the product |
| | Product Edition | Product edition defined by the vendor. For example, Professional |
| | Path | Directory path where the product is installed on the computer system |
| | AntiVirus Definition Date | The date of the virus definition file installed on the computer. Some anti-virus products combine date and version into a single string. NOTE: This is applicable only for antivirus products. |
| | AntiVirus Definition Version | The vendor-defined version of the virus definition file that has been installed on a computer. NOTE: This is applicable only for antivirus products. |
|  Hardware/Software Inventory >  Software >  Application Vendors > <i>Vendor_name</i> >  <i>software_group_name</i> >  <i>software</i> >  Patches | Name | Vendor-defined name for the patch |
|  Hardware/Software Inventory >  Software >  Application Vendors > <i>Vendor_name</i> >  <i>software_group_name</i> >  <i>software</i> >  Representative File Information | File Name | Name of the file representing the software |
| | File Version | Version of the file representing the software |
| | File Size | Size of the file representing the software |
| | Last Modified | Last modified date of the file representing the software |
| | Internal Name | Internal name |
| | Product Version | The version of the product represented by this file |
| | Company | Vendor name |
| | Product Name | The product that this file represents |
| | Language | User-friendly name for the language of this copy of the file |
| | File Path | Location of the file on the inventoried server |
| | Software Dictionary ID | ID of the file as represented in the General software dictionary |

| Scan Data Group | Scan Data Item | Description |
|---|-----------------------------|--|
|  Hardware/Software Inventory >  Software >  Disk Usage | File Extension Name | The file extension for which the disk usage is scanned for. |
| | Total Disk Usage | Total disk usage for all the files of the specified extension. |
|  Hardware/Software Inventory >  Software >  Device Drivers >  Pointing Device Drivers >  <i>Pointing Device driver name</i> | Name | Name of the mouse driver |
| | Version | Version number of the mouse driver |
|  Hardware/Software Inventory >  Software >  Device Drivers >  Display Drivers | Install Date | Install date of the display driver |
| | Manufacturer | Name of the display driver manufacturer |
| | Is Shadowed (True or False) | If True, the display driver is currently being shadowed |
| | Version | Version number of the display driver |
|  Hardware/Software Inventory >  Software >  Device Drivers >  Network Drivers | Description | Description of the network driver |
| | Name | Network driver name |
| | Version | Version number of the network driver |
|  Hardware/Software Inventory >  Software >  Operating System | Code Page | Language code page of the operating system |
| | OS Type | Operating system of the inventoried server |
| | Install Date | Install date of the operating system |
| | Caption | Operating system name, for example, Windows 95/Windows 2000 |
| | Other Description | Additional description of the operating system if available |
| | Role | Type of the operating system such as server or workstation |
| | Total Virtual Memory Size | Total number of bytes in the virtual address space of the calling process |
| | Total Memory Size | Total memory of the operating system |
| | Version | Version number of the operating system |
|  Hardware/Software Inventory >  Hardware >  Monitor | Device ID | Unique ID of a desktop monitor that is attached to a computer system For example, DesktopMonitor1 |

| Scan Data Group | Scan Data Item | Description |
|---|-----------------------|--|
| | Description | Description of the monitor. |
| | Nominal Size | <p>A number representing the diagonal width of the monitor (the distance from one corner of the screen to the opposite corner of the screen)</p> <p>For example, 17"</p> <p>You can customize the scan of the nominal size of the monitor by configuring the HWRules ini file using the Server Inventory policy.</p> |
| | Viewable Size | <p>A number representing the diagonal width of the screen image excluding the black borders around the image's edge</p> <p>For example, 15.8"</p> |
| | Manufacturer | <p>Name of the monitor's manufacturer</p> <p>For example, DELL* Computer Corp</p> |
| | Serial Number | <p>Manufacturer's number used to identify a monitor</p> <p>For example, 23DDC24N9067</p> |
| | Model | <p>Product name of the monitor given by the manufacturer</p> <p>For example, DELL E771a</p> |
| | Manufacture Date | <p>Year in which the monitor was manufactured</p> <p>For example, 2003</p> |
| | Model ID | <p>Unique ID of a model of the monitor; it is a combination of the Manufacturer ID and Product ID</p> <p>For example, DELA001</p> |
|  Hardware/Software Inventory >  Hardware >  Chassis | Asset Tag | <p>Asset tag number of the system chassis</p> <p>For example, S11127</p> |
| | Number of Power Cords | Total number of power cords attached to a system chassis |
| | Chassis Type | Represents whether the system chassis is a laptop, desktop, notebook, docking station and so on |
| | Manufacturer | <p>Name of the system chassis manufacturer</p> <p>For example, Compaq</p> |

| Scan Data Group | Scan Data Item | Description |
|---|-------------------------------|---|
| | Serial Number | Manufacturer's number used to identify a system chassis For example, 53R661S |
| | Version | Version number of the system chassis |
| | Tag | Unique ID of the system chassis attached to a particular computer system For example, System Enclosure 0 |
|  Hardware/Software Inventory >  Hardware >  Pointing Device >  <i>Pointing device name</i> | IRQ Number | Interrupt assigned to this device |
| | Name | Identifying information of the mouse |
| | Number of Buttons | Number of buttons on the mouse |
|  Hardware/Software Inventory >  Hardware >  Keyboard | Delay | Delay before the repeat of a key |
| | Description | Description of the keyboard, such as IBM Enhanced 101 or 102 keys |
| | Layout | Layout of the keyboard |
| | Number of Function Keys | Total number of function keys |
| | Subtype | Type of the keyboard |
| | Typematic Rate | Rate of processing the keys |
|  Hardware/Software Inventory >  Hardware >  Display Adapter >  <i>Display adapter name</i> | Chip Set | Chip set used by the controller to compare system capabilities |
| | Current Bits/Pixel | Number of adjacent color bits for each pixel |
| | Current Horizontal Resolution | Number of horizontal pixels shown by the display |
| | Current Vertical Resolution | Number of vertical pixels shown by the display |
| | DAC Type | Digital-to-Analog converter type |
| | Description | Description of the display adapter |
| | Maximum Memory Supported | Maximum memory that the display adapter supports for VIDEO RAM |
| | Maximum Refresh Rate | Maximum refresh rate of the monitor for redrawing the display, measured in Hertz |

| Scan Data Group | Scan Data Item | Description |
|---|------------------------------|--|
| | Minimum Refresh Rate | Minimum refresh rate of the monitor for redrawing the display, measured in Hertz |
| | Number of Color Planes | Number of color planes supported by the video system |
| | Provider | Vendor name |
| | Video Architecture | The architecture of the video subsystem in this system, for example, CGA/VGA/SVGA/8514A |
| | Video Memory Type | The type of video memory for this adapter, for example, VRAM/SRAM/DRAM/EDO RAM |
|  Hardware/Software Inventory >  Hardware >  BIOS | BIOS Identification Bytes | Byte in the BIOS that indicates the computer model |
| | Install Date | The manufacturing date of the BIOS |
| | Manufacturer | BIOS vendor name |
| | Caption | BIOS label |
| | Primary BIOS | True state indicates Primary BIOS |
| | Serial Number | Serial number of the computer, assigned during manufacture |
| | Size | Size of the BIOS |
| | Version | Version or revision level of the BIOS |
|  Hardware/Software Inventory >  Hardware >  Processor | Current Clock Speed (in MHz) | Current clock speed of the processor |
| | Device ID | Special hexadecimal string identifying the processor type |
| | Maximum Clock Speed (in MHz) | Maximum clock speed of the processor |
| | Other Family Description | Additional description about the Processor Family, such as Pentium Processor with MMX technology |
| | Processor Family | Identification of the processor family such as Pentium II or Pentium III |
| | Processor Stepping | Single-byte code characteristic provided by microprocessor vendors to identify the processor model |
| | Role | Type of processor such as central processor, math coprocessor, and others |
| | Upgrade Method | The method by which this processor can be upgraded, if upgrades are supported |

| Scan Data Group | Scan Data Item | Description |
|---|------------------------------------|--|
|  Hardware/Software Inventory >  Hardware >  Modem | Description | Additional information about the modem |
| | Name | Identifying information of the modem |
| | Device ID | Special hexadecimal string identifying the modem type |
| | Provider | Name of the vendor |
|  Hardware/Software Inventory >  Hardware >  Battery | Chemistry | The battery chemistry, for example, lithium-ion or nickel metal hydride |
| | Design Capacity | The design capacity of the battery in mWatt-hours |
| | Design Voltage | The design voltage of the battery in mVolts |
| | Install Date | The battery manufacture date |
| | Manufacturer | The name of the company that manufactured the battery |
| | Name | Device name for this battery, for example, Duracell* DR-36 |
| | Serial Number | The serial number for this battery |
|  Hardware/Software Inventory >  Hardware >  Power Supply | Description | Expanded description of the input voltage capability for this power supply |
| | Total Output Power (in MilliWatts) | Attribute value that represents the total output power of the power supply |
|  Hardware/Software Inventory >  Hardware >  Memory | Total Memory | Total memory of the inventoried server |
|  Hardware/Software Inventory >  Hardware >  Disk Drives >  Floppy | Capacity | Floppy drive capacity |
| | Description | Floppy drive description |
| | Drive Letter | Letter name of the drive |
| | Manufacturer | Vendor name |
| | Physical Cylinders | Floppy drive cylinders |
| | Physical Heads | Floppy drive R/W heads |
| Sectors/Track | Floppy drive sectors per track | |
|  Hardware/Software Inventory >  Hardware >  Disk Drives >  Physical Disk >  Fixed Disk | Description | Description |

| Scan Data Group | Scan Data Item | Description |
|---|----------------------|--|
| | Manufacturer | Vendor name |
| | Physical Cylinders | Number of cylinders |
| | Physical Heads | Number of heads |
| | Sectors/Track | Fixed disk drive sectors per track |
| | Size | Size of the fixed disk |
|  Hardware/Software Inventory >  Hardware >  Disk Drives >  Physical Disk >  Removable Disk | Description | Description |
| | Manufacturer | Vendor name |
| | Physical Cylinders | Number of cylinders |
| | Physical Heads | Number of heads |
| | Sectors/Track | Removable disk drive sectors per track |
| | Size | Size of the removable disk |
|  Hardware/Software Inventory >  Hardware >  Disk Drives >  Logical Disk >  Logical disk name | Drive Letter | Letter name of the drive |
| | File System Type | Type of file system, such as File Allocation Table (FAT) |
| | Free Size | Drive's actual size in MB |
| | Volume Label | Name of the hard disk volume |
| | Size | Drive's available space in MB |
| | Volume Serial Number | Hard disk volume serial number |
|  Hardware/Software Inventory >  Hardware >  Disk Drives >  CDROM | Name | Name of the CD drive attached to the inventoried server |
| | Description | Description of the CD |
| | Drive Letter | Mapped drive name of the CD |
| | Manufacturer | Vendor name |
| | Caption | Caption of the CD |
|  Hardware/Software Inventory >  Hardware >  Ports >  Serial Port | Address | Base input/output address for this serial port |
| | IRQ Number | IRQ number of the serial port |
| | Name | The logical name of the I/O device on this serial port, under this operating environment |

| Scan Data Group | Scan Data Item | Description |
|---|-----------------------------|---|
|  Hardware/Software Inventory >  Hardware >  Ports >  Parallel Port | Address | Base I/O address for this parallel port |
| | DMA Support (True or False) | If True, DMA is supported |
| | Name | The logical name of the input-output device on this parallel port, under this operating environment |
| | IRQ Number | IRQ number of the parallel port |
|  Hardware/Software Inventory >  Hardware >  Bus | Bus Type | Bus type indicates PCI, ISA, and others |
| | Description | Bus description |
| | Name | Bus name |
| | Version | Version of the bus supported by the motherboard |
|  Hardware/Software Inventory >  Hardware >  Network Adapter | Adapter Type | Type of network adapter, such as FDDI or token ring |
| | Auto Sense | A Boolean value indicating whether the network adapter is capable of automatically determining the speed or other communication characteristics of the attached network media |
| | Card Manufacturer | Name of the card manufacturer |
| | Description | Adapter description |
| | Install Date | Install date of the network adapter |
| | Maximum Speed | Rate at which the information is transferred over the LAN |
| | Name | Network adapter name |
| | Permanent Address | Node address stored permanently in the adapter |
| | Provider | Name of the provider |
|  Hardware/Software Inventory >  Hardware >  Sound Adapter | Description | Description of the multimedia component for the server |
| | Name | Label of the multimedia card |
| | Provider | Name of the provider |
|  Hardware/Software Inventory >  Network >  DNS | DNS Name | The DNS name of the inventoried server |

| Scan Data Group | Scan Data Item | Description |
|---|-------------------|--|
|  Hardware/Software Inventory >  Network >  Network (instance_number) >  IP | IP Address | The unique address assigned to a computer on an IP Internet |
| | Subnet Mask | The subnet mask of the inventoried server paired with an IP address specifies to an IP router which octets or bits in the IP address are the network ID and which octets or bits are the node ID |
|  Hardware/Software Inventory >  Network >  Network (instance_number) >  IPX | IPX Address | The IPX™ address of the inventoried server |
|  Hardware/Software Inventory >  Network >  Network (instance_number) >  MAC | MAC Address | Unique node address permanently coded in the network adapter that identifies a specific computer on a network |
|  Hardware/Software Inventory >  Network > IP | IP Address | The unique address assigned to a computer on an IP Internet |
| | Subnet Mask | The subnet mask of the inventoried server paired with an IP address specifies to an IP router which octets or bits in the IP address are the network ID and which octets or bits are the node ID |
| Hardware/Software Inventory > Network > IPX | IPX Address | The IPX address of the inventoried server |
| Hardware/Software Inventory > Network > MAC | MAC Address | Unique node address permanently coded in the network adapter that identifies a specific computer on a network |
|  Hardware/Software Inventory >  System >  System IRQ | Availability | Availability of the specific IRQ channel |
| | IRQ Number | Number of the Interrupt Request Line (IRQ), from 0 to 15 |
| | IRQ Trigger Type | IRQ Trigger type |
|  Hardware/Software Inventory >  System >  System Cache | Shareable | If True, the system IRQ can be shared across devices |
| | Associativity | Defines the system cache associativity (direct-mapped, 2-way, 4-way) |
| | Cache Type | Defines the system cache type, for example, Instruction, Data, Unified |
| | Capacity | Size of the data store where the cache information is kept |
| | Error Methodology | Error correction scheme supported by this cache component, for example, Parity/Single Bit ECC/MultiBit ECC |

| Scan Data Group | Scan Data Item | Description |
|---|--------------------|---|
| | Level | Indicates the cache level; internal cache that is built in to the microprocessors; external cache that is between the CPU and DRAM |
| | Line Size | Size in bytes of a single cache bucket or line |
| | Read Policy | Indicates whether the data cache is for read operations |
| | Replacement Policy | Algorithm that the cache uses to determine which cache lines or buckets should be reused |
| | Speed | Speed of this System Cache module in nanoseconds |
| | Write Policy | Indicates the two different ways (Write-Back and Write-Through Cache) that the cache can handle to write to the memory |
|  Hardware/Software Inventory >  System >  System DMA | Availability | Indicates whether Virtual Direct Memory Access (DMA) is supported |
| | Description | Name of the logical device that is currently using this DMA channel |
| | DMA Burst Mode | A data transmission mode in which data is sent faster than normal |
| | DMA Channel Number | Number of the Direct Memory Access (DMA) channel that a computer uses for transferring data to and from devices quicker than from computers without a DMA channel |
|  Hardware/Software Inventory >  System >  System Slot | Description | Card currently occupying this slot |
| | Maximum Data Width | Maximum bus width of cards accepted in the slot |
| | Thermal Rating | Maximum thermal dissipation of the slot in milliwatts |
|  Hardware/Software Inventory >  System >  Motherboard | Manufacturer | Name of the motherboard manufacturer |
| | Number of Slots | The number of expansion slots in the motherboard for adding more memory, graphic capabilities, and support for special devices |
| | Version | Version of the motherboard |
| | Description | General description of the motherboard |

NOTE: For an enumerated attribute, the value is displayed in the format *enumerated_value* [*enumerated_ID*]. For example, Processor.Processor Family = Pentium (R) III [17].

The Status bar displays the following information:

- ♦ **Tree Name:** Displays the eDirectory tree name where the inventoried server or inventoried server resides.
- ♦ **Recent Information:** Set to *Yes* if the Inventory database has been updated with the latest inventory information of the selected inventoried server.

17.1.3 Viewing Inventory Information of Inventoried Servers by Querying the Database

Using ConsoleOne, you can query the Inventory database to display the hardware and software components of inventoried servers that you want to view. The Inventory Query window displays the information satisfying the criteria you specify.

The Inventory database stores inventory information (general, hardware, software, network, and system information) for each inventoried server. Querying the Inventory database helps to create groups of similar devices and to focus your reports on specific types of machines. For example, you can query the database to find machines that have an i486D processor and a VGA card.

To query the Inventory database for inventory information:

1 In ConsoleOne, click a container.

2 Invoke Query:

- ♦ To invoke the Inventory query from a database object, right-click the database object, click *ZENworks Inventory*, then click *Inventory Query*.
- ♦ To invoke the Inventory query from the ConsoleOne Tools menu, you must first configure the database and then click *Tools*, click *ZENworks Inventory*, then click *Inventory Query*. For more information on how to configure the Inventory database, see [“Configuring the Inventory Database” on page 635](#).

3 Specify the criteria for query:

Query the Inventory database for: By default, the *Servers* option is enabled. The query locates all inventoried servers satisfying the query expression. If ZENworks 7 Server Management and Desktop Management are installed in the same environment; the *Workstations*, the *Servers* and the *Both* options are available. When you select *Servers*, the query locates all inventoried servers satisfying the query expression. Choose *Both* to include all workstations and inventoried servers satisfying the query expression.

Find Type: Select *Quick* or *Advanced*. Click *Quick* to specify a simple query. When you choose a *Quick* query, you specify one attribute, relational operators, and the value of the attribute. Choose *Advanced* query to specify many attributes. You can combine multiple query groups so each group defines a set of query criteria. For example, use the Advanced query to run a query to discover all devices in the database with 486 processors and use query connectors, and add another query to discover which of these inventoried servers have a VGA color video adapter.

Display Machine(s) Not Satisfying the Query: Select the check box to retrieve machines that do not satisfy the query.

Select Attribute: Select the component or component attributes. Attributes that you can specify to query on the inventoried servers are grouped into the following categories: General, Software, Hardware, Network, and System.

The custom attribute is prefixed by an asterisk (*).

For example, to find the machines that do not have a pointing device installed, select Pointing Device as the component. To specify the version of BIOS as a component in the query, select BIOS as the component and VERSION as the component attribute.

Operator or Relational Operator: Select to determine the relationship between the components and the value. The relational operators are grouped on the basis on the data type of the attribute selected in the Select Attribute window as shown in the following table:

| Data Type of the Attribute | Relational Operators |
|----------------------------|---|
| String | Equal To (=), Not Equal To (!=), Matches ([]), Does Not Match (![]) and Is NULL (null) |
| Numeric | Equal (=), Not Equal (!=), Less Than (<), Less Than or Equal To (<=), Greater Than (>), Greater Than or Equal To (>=), and Is NULL (null) |
| Date | After (>), On or After (>=), Before (<), On or Before (<=), and Is NULL (null) |
| Enum | Equal To (=), Not Equal To (!=), and Is NULL (null) |
| Custom | Includes all the relational operators that are grouped under the String, Numeric, and Date data types |

NOTE: If the query does not display the result when the data type of the attribute is Custom and the relational operator is Numeric or Date, use the Equal To operator to find the values for the custom attributes that are stored in the Inventory database.

If you select only the component in the Select Attribute window, the Relational Operator is set to NULL by default and other relational operators are not available.

Value: Description values are the possible values of an inventory component. For example, 6.0 is a possible value for the DOS-Version attribute. Description values are not case sensitive.

NOTE: For an enumerated attribute, the value is displayed in the format, *enumerated_value [enumerated_ID]*. For example, Processor.Processor Family = Pentium (R) III [17].

If you choose Matches ([]) or Does Not Match (![]) as the relational operator, you can use wildcards to substitute characters in the Value field. The following table lists the wildcards that can be used according to the SQL documentation:

| Example | Specifies to Include |
|----------------|---|
| ? | Any one character |
| _ (underscore) | Any one character |
| % | Any string of zero or more characters |
| [] | Any one character in the specified range or set |
| [^] | Any one character not in the specified range or set |

NOTE: To define a query using special characters such as ? or [, specify the query in the following formats: [?] or [[]].

The list of description values displayed for an Inventory component is taken from the Inventory database corresponding to the component.

Logical Operator: This option is available only for the Advanced query. Logical Operator forms query groups that is combined with the previous query group by using the relational operator specified between the query groups.

Save: This option is available only for the Advanced query. It saves the query expression as a file in the location that you specify. The query file does not have a default extension; however, we recommend the `.qry` extension for easy reference.

Load: This option is available only for the Advanced query. It loads the query file that you specify. You must provide the full filename with its extension.

4 Click *Find*.

This will query based on the query criteria you specify and display the inventoried servers that match the query in the Query Results window.

In the Query Results window, double-click the inventoried server or click *File*, then click *Advanced Query* to view the **inventory information** of the inventoried server.

Usage of Relational Operators

- ◆ **Match:** Use the Match operator to find the inventoried servers that satisfy the query condition.

For example, use the Match operator to find all the inventoried servers with IP address 164.99.151.%.

- ◆ **NULL:** Use the NULL operator to query for those inventoried servers whose particular attribute is not scanned but the component has been scanned and some attributes are populated.

For example, to find a list of inventoried servers for which BIOS.Manufacturer is not scanned, form a BIOS.Manufacturer is NULL query. This query displays the inventoried servers for which the BIOS has been scanned.

- ◆ **NOT SATISFYING:** Use the NOT SATISFYING query (or the NOT SATISFYING filter condition) to find filter conditions for the inventoried servers that negate the given query.

For example, two servers S1 and S2 contain serial ports COM1 and COM2. The query (SerialPort='COM1') returns S1 and the query (SerialPort!='COM1') also returns the S1 because S1 contains the serial port COM2. To query the inventoried servers that do not contain the serial port COM1 you must use <NOT SATISFYING>(SerialPort='COM1'). To use the NOT SATISFYING option, click the *Display Machines Not Satisfying the Query* check box in the query window.

17.1.4 Running Inventory Reports

You can run reports to gather inventory information from the Inventory database. The Inventory reports are designed using Crystal Reports*.

You can select from a predefined set of report forms to generate a report. The inventory report is displayed in the Crystal Viewer window.

You can print or export the report as desired. Remember that any reports you generate would be empty if you have not configured ZENworks 7 Server Management to start populating the Inventory database with the information you want.

This section covers information on the following sections:

- ◆ “Prerequisites for Generating Inventory Reports” on page 652
- ◆ “Types of Inventory Reports” on page 652
- ◆ “Generating Inventory Reports” on page 660
- ◆ “Printing an Inventory Report” on page 661
- ◆ “Exporting an Inventory Report to a File” on page 662
- ◆ “Understanding User-Defined Reports” on page 662

Prerequisites for Generating Inventory Reports

Before running the inventory reports, ensure that you have installed the appropriate ODBC client. For more information, see “Installing the ODBC Drivers” in the *Novell ZENworks 7 Server Management Installation Guide*.

Types of Inventory Reports

You can generate the types of reports described below, assuming you have already configured ZENworks 7 Server Management to start populating the inventory database with the information you want.

Table 17-2 gives the Simple Inventory lists that provide information on individual aspects of Server Inventory, such as the operating system and the selection criteria. The table also lists the Comprehensive Inventory Reports that combine several aspects of Server Inventory into each report, such as memory, hard disk, and processor.

Table 17-2 List of Inventory reports and information displayed by each report

| Inventory Report Group | Report Name | Selection Criteria | Information Displayed in the Inventory Report |
|------------------------|-------------------------|--|--|
| Hardware Inventory | Asset Management Report | Scope, Machine Name, IP Address, and DNS Name You can also select to display the following options in the report: Memory, Processor, Display Adapter, Keyboard, Pointing Device, Fixed and Removable Disk, Floppy, CDROM, Network Adapter, and Monitor | Memory, processor, display details, keyboard, pointing device, fixed and removable disk, floppy, CD drive, network adapter, and monitor details for inventoried servers. |
| | BIOS Listing | Scope, Machine Name, IP Address, DNS Name, BIOS Install Date, and Manufacturer | List of all the inventoried servers with BIOS manufacturer, BIOS release date, and the total number of such machines. |
| | Battery Listing | Scope, Machine Name, IP Address, DNS Name, and Name | List of all inventoried servers that match the specified battery name. |
| | Bus Listing | Scope, Machine Name, IP Address, DNS Name, and Bus Type | List of all inventoried servers with the selected bus type. |

| Inventory Report Group | Report Name | Selection Criteria | Information Displayed in the Inventory Report |
|-------------------------------|-------------------------|--|--|
| | CDROM Listing | Scope, Machine Name, IP Address, DNS Name, Caption, Description, and Manufacturer | List of all inventoried servers that match the specified CD caption, description, and manufacturer's name. |
| | Display Adapter Listing | Scope, Machine Name, IP Address, DNS Name, Video Architecture, and Description | List of all inventoried servers that match the specified display adapter's video architecture and description. |
| | Floppy Listing | Scope, Machine Name, IP Address, DNS Name, Manufacturer, and Description | List of all inventoried servers that match the specified floppy description and manufacturer's name. |
| | Hardware Summary Report | Scope, Machine Name, IP Address, DNS Name, Operating System Type, Operating System Version, Processor Family, Curr. Clock Speed (Lower Bound in MHz), Curr. Clock Speed (Upper Bound in MHz), Total Memory (Lower Bound in MB), Total Memory (Upper Bound in MB), Hard Disk Size (Lower Bound in GB), and Hard Disk Size (Upper Bound in GB) | Operating system name, operating system version, processor family, processor current clock speed, memory, and hard disk size for each inventoried server. |
| | Keyboard Listing | Scope, Machine Name, IP Address, DNS Name, Description, and Layout | List of all inventoried servers that match the specified keyboard description and layout. |
| | Modem Listing | Scope, Machine Name, IP Address, DNS Name, and Name | List of all inventoried servers that match the specified modem name. |
| | Monitor Listing | Scope, Machine Name, IP Address, DNS Name, Manufacturer, Manufacture Date, Nominal Size (Lower Bound in inches), and Nominal Size (Upper Bound in inches) | List of all inventoried servers that match the specified monitor manufacturer's name, manufacture date, and the specified range of monitor's nominal size. |
| | Network Adapter Listing | Scope, Machine Name, IP Address, DNS Name, and Name | List of all inventoried servers that match the specified network adapter's name. |
| | Physical Disk Listing | Show Chart, Scope, Machine Name, IP Address, DNS Name, Removable, Manufacturer, Description, Total Size (Lower Bound in GB), and Total Size (Upper Bound in GB) | List of all inventoried servers that match the specified physical disk manufacturer's name, description, the specified range of total size and disks that are fixed, removable, or both. You can also select the Show Chart box to display the Physical Disk Listing report in a pie chart. |
| | Pointing Device Listing | Scope, Machine Name, IP Address, DNS Name, Pointing Device Type, and Pointing Device Name | List of all inventoried servers that match the specified pointing device type and name. |

| Inventory Report Group | Report Name | Selection Criteria | Information Displayed in the Inventory Report |
|-----------------------------------|----------------------------------|---|--|
| | Power Supply Listing | Scope, Machine Name, IP Address, DNS Name, and Description | List of all inventoried servers that match the specified power supply description. |
| | Processor Listing | Show Chart, Scope, Machine Name, IP Address, DNS Name, Processor Family, Maximum Speed (Lower Bound in MHz), Maximum Speed (Upper Bound in MHz), Current Speed (Lower Bound in MHz), and Current Speed (Upper Bound in MHz) | List of all the inventoried servers with a processor family (such as Pentium Pro), processor maximum clock speed, and the processor current clock speed of the machines. You can also select the Show Chart box to display the Processor Listing report in a pie chart. |
| | Sound Adapter Listing | Scope, Machine Name, IP Address, DNS Name, and Name | List of all inventoried servers that match the specified sound adapter name. |
| | Storage Devices Inventory Report | Scope, Machine Name, IP Address, and DNS Name You can also select to display the following options in the report: Fixed and Removable Disk, Logical Disk, Floppy, and CDROM. | Fixed disk, removable disk, logical disk, floppy, and CD drive details for each inventoried server. |
| | System Chassis Listing | Scope, Machine Name, IP Address, DNS Name, Chassis Type, and Manufacturer | List of all inventoried servers that match the specified system chassis type and manufacturer's name. |
| System Configuration on Inventory | Inventory Scan Listing | Show Chart, Scope, Machine Name, IP Address, DNS Name, Last Scan Date (On or Before), Inventory Server Name, and Recent Information | Date and time of the last inventory scan, Inventory server name, and recent information on each inventoried server. You can also select the Show Chart box to display the System Configuration Inventory report in a pie chart. |
| | Memory Listing | Show Chart, Scope, Machine Name, IP Address, DNS Name, Total Memory (Lower Bound in MB), and Total Memory (Upper Bound in MB) | List of all the inventoried servers within a range of memory size (such as 200-400 MB) and the total number of such machines. You can also select the Show Chart box to display the Memory Listing report in a pie chart. |
| | Operating System Listing | Show Chart, Scope, Machine Name, IP Address, DNS Name, Operating System Type, and Operating System Version | List of all inventoried servers that match the specified operating system type and version. You can also select the Show Chart box to display the Operating System Listing in a pie chart. |

| Inventory Report Group | Report Name | Selection Criteria | Information Displayed in the Inventory Report |
|-------------------------------|---|---|--|
| | Networking Information Report | Scope, Machine Name, IP Address, and DNS Name | Network adapter type, DNS, IP address, MAC address, IPX address, and Windows Domain name for each inventoried server. |
| | System Information Listing | Scope, Machine Name, IP Address, DNS Name, and Computer Manufacturer | List of all inventories servers that match the specified computer manufacturer's name. |
| | System Internal Hardware Inventory Report | Scope, Machine Name, IP Address, and DNS Name You can also select to display the following options in the report: System IRQ, System Cache, System DMA, System Slot, and Motherboard. | IRQ, cache, DMA, slot, and motherboard for each inventoried server. |
| Software Inventory | Add-Remove Programs by Application | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, and Software Name | List of all software that are listed in the Add/Remove Programs list for each inventoried server. |
| | Add-Remove Programs by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, and Software Name | List of all software that are listed in the "Add-Remove Programs" list for each inventoried server. |
| | Anti-Virus Signature Files by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, Software Name, Min Virus Definition Date and Max Virus Definition Date | List of all antivirus signature files grouped by antivirus product installed on each inventoried server. |
| | Anti-Virus Signature Machine Count | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, Software Name, Min Virus Definition Date and Max Virus Definition Date | List showing the count of inventoried servers that have any antivirus product installed. |
| | Disk Usage by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, and File Extension | List of all inventoried servers and the disk usage that match the specified file extension. |
| | Exception List by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, File Name, Vendor Name, and Product Name | List of all inventoried servers and the file information that match the specified filename, vendor name, and product name. |
| | Installed NetWare Software by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, and Software Name | List of all inventoried NetWare machines and the products.dat details that match the given software name. |

| Inventory Report Group | Report Name | Selection Criteria | Information Displayed in the Inventory Report |
|-------------------------------|--|--|--|
| | Internet Explorer Installation Count | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, Internet Explorer Version, and Service Pack | List showing the count of inventoried servers with Internet Explorer installed. |
| | Internet Explorer Patches by Machine | Scope, Machine Name, IP Address, DNS Name, Internet Explorer Version, and Service Pack | List of all installed patches for the Internet Explorer version that matches the specified value and patch name. |
| | Internet Explorer by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, Internet Explorer Version, and Service Pack | List of all Internet Explorer installations that match the specified version. |
| | MSI Products by Application | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, Software Name, and Vendor Name | List of all products installed on each inventoried server and that are listed in the MSI (Microsoft Installer) database. |
| | MSI Products by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, Software Name, and Vendor Name | List of all products installed on each inventoried server and that are listed in the MSI (Microsoft Installer) database. |
| | Microsoft Office Components by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, Microsoft Office Version, and Service Pack | List of all products that match the specified product name and vendor name, and have been installed from the specified source. |
| | Microsoft Office Installation Count | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, Microsoft Office Version, and Service Pack | List showing the count of inventoried servers with Microsoft Office installed. |
| | Microsoft Office by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, Microsoft Office Version, and Service Pack | List of all Microsoft Office installations that match the specified version. |
| | Novell Client Components by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, Novell Client Version, and Service Pack | List of all Novell Client components that match the specified version. |
| | Novell Client Installation Count | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, Novell Client Version, and Service Pack | List showing the count of inventoried servers with Novell Client installed. |
| | Novell Client by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, Novell Client Version, and Service Pack | List of all Novell Client installations that match the specified version. |

| Inventory Report Group | Report Name | Selection Criteria | Information Displayed in the Inventory Report |
|-------------------------------|--|---|--|
| | Outlook Express Installation Count | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, and Outlook Express Version | List showing the count of inventoried servers with Outlook Express installed. |
| | Outlook Express by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, and Patch Name | List of all Outlook Express installations that match the specified version. |
| | Software Dictionary Application Files by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, Vendor Name, Software Name, and Software Version | List of all inventoried servers and their software dictionary application files that match the specified vendor, software, and software version. |
| | Software Dictionary Applications by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, Vendor Name, Software Name, and Software Category | List of all inventoried servers and their software dictionary applications that match the specified vendor, software, and software version. |
| | Software Dictionary Versions Machine Count | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, General Dictionary Version, and Private Dictionary Version | List showing the count of all inventoried servers with specified General Dictionary and Private Dictionary versions. |
| | Software Dictionary Versions by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, General Dictionary Version, and Private Dictionary Version | List of all inventoried servers with specified General Dictionary and Private Dictionary versions. |
| | Software Installation Count | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, Vendor Name, Software Name, and Software Version | List showing the count of inventoried servers with specified vendor name, software, and version. |
| | Software Installations | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, Vendor Name, Software Name, and Software Version | List of all inventoried servers with specified vendor name, software, and version. |
| | Software by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, Vendor Name, Software Name, and Software Version | List of all inventoried servers and software information that match the specified vendor name, software, and version. |

| Inventory Report Group | Report Name | Selection Criteria | Information Displayed in the Inventory Report |
|-------------------------------|--|--|---|
| | System Software Inventory Report | Scope, Machine Name, IP Address, and DNS Name. You can also select to display the following options in the report: Display Driver, Pointing Device Driver, Network Adapter Driver, and NetWare Client. | Drivers (such as pointing device drivers, network adapter drivers, and display drivers) and Novell NetWare® Client for each inventoried server. |
| | Windows Components by Machine | Scope, Machine Name, IP Address, DNS Name, Windows Version, and Service Pack | List of all Windows components that match the specified version. |
| | Windows Installation Count | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, Windows Version, and Service Pack | List showing the count of inventoried servers that have Windows operating system installed. |
| | Windows Media Player Count | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, and Windows Media Player Version | List showing the count of inventoried servers with Windows Media Player installed. |
| | Windows Media Player Patches by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, and Windows Media Player Version | List of all patches for Windows Media Player installations that match the specified version and patch name. |
| | Windows Security Patches by Patch | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, Windows Version, and Service Pack | List of all patches for Windows operating systems that match the specified version and patch name. |
| | Windows Media Player by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, and Windows Media Player Version | List of all Windows Media Player installations that match the specified version. |
| | Windows Operating System by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, Windows Version, and Service Pack | List of all Windows operating systems that match the specified version and serial number. |
| | Windows Security Patches by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, Max Last Scan Time, Windows Version, and Service Pack | List of all patches for Windows operating systems that match the specified version and patch name. |
| | Novell ZENworks Desktop Management Agent Components by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, and Max Last Scan Time | List of all machines that were successfully last scanned within the specified time range and the agent components of ZENworks 7 Desktop Management installed on these machines. |

| Inventory Report Group | Report Name | Selection Criteria | Information Displayed in the Inventory Report |
|-------------------------------|---|---|--|
| | Novell ZENworks Desktop Management Server Components by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, and Max Last Scan Time | List of all machines that were successfully last scanned within the specified time range and the server components of ZENworks 7 Desktop Management installed on these machines. |
| | Novell ZENworks Handheld Management Components by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, and Max Last Scan Time | List of all machines that were successfully last scanned within the specified time range and the ZENworks 7 Handheld Management components installed on these machines. |
| | Novell ZENworks Installed Components by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, and Max Last Scan Time | List of all machines that were successfully last scanned within the specified time range and the ZENworks 7 components installed on these machines. |
| | Novell ZENworks Installed Suites by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, and Max Last Scan Time | List of all machines that were successfully last scanned within the specified time range and the ZENworks 7 suites installed on these machines. |
| | Novell ZENworks Server Management Agent Components by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, and Max Last Scan Time | List of all machines that were successfully last scanned within the specified time range and the agent components of ZENworks 7 Server Management installed on these machines. |
| | Novell ZENworks Server Management Server Components by Machine | Scope, Machine Name, IP Address, DNS Name, Min Last Scan Time, and Max Last Scan Time | List of all machines that were successfully last scanned within the specified time range and the server components of ZENworks 7 Server Management installed on these machines. |

| Inventory Report Group | Report Name | Selection Criteria | Information Displayed in the Inventory Report |
|------------------------|---|---|---|
| Others | User Defined Reports | Based on the options specified by the user in the <code>consoleone\consoleone_version\bin\userreports.ini</code> file | Displays the user-defined report. |
| | For more information on how to create user-defined reports, see the “Understanding User-Defined Reports” on page 662. | | |

NOTE: The Show Chart selection criteria display a graphical representation of the Inventory report.

Generating Inventory Reports

To generate the inventory report:

- 1 Invoke the Inventory report by using any of the following methods:
 - ♦ To invoke the Inventory report from a database object, right-click the database object, then click *ZENworks Reports*.
 - ♦ To invoke the Inventory report from the ConsoleOne Tools menu, you must first configure the database, click *Tools*, then click *ZENworks Reports*. For more information on how to configure the database, see [“Configuring the Inventory Database”](#) on page 635.

- 2 Click the report you want to generate.

The description for the report is displayed on the right side of the screen.

See the table with listing of simple Inventory lists and listing of the comprehensive inventory reports.

- 3 Specify the selection criteria.

The Scope selection criteria is enabled only if both ZENworks 7 Desktop Management and ZENworks 7 Server Management are installed on the same machine.

For example, to view all inventoried servers that have the Windows 2000 operating system, you would select Operating System Listing, specify the selection criteria scope as *Both*, and the operating system type as Windows 2000. The report displays the inventory information of all servers within the configured Inventory database.

Depending on the type of report you want, you can filter the information. For example, to view all inventoried servers with the Windows operating system, you select the Operating System Listing, and specify the selection criteria Scope as *Both*, the Operating System Type as Windows, and the Operating System Version as 2000.

Follow these guidelines as you work with the Reporting dialog box:

- ♦ The selection criteria in the Inventory report are case sensitive

For example, if you want to know the list of machines whose Distinguished Name is CN=MACHINE1.OU=ENG.O=NOVELL, specify OU=ENG.O=NOVELL as the selection criterion. All the machines whose DN contains OU=ENG.O=NOVELL are displayed in the Inventory report, but the machines whose DN contains ou=eng.o=novell are not displayed in the Inventory report.

- ◆ If the Reporting dialog box allows wildcards, you can use an asterisk (*) or question mark (?) with all selection criteria. The wildcard characters can be used for text fields only.

You can use * to retrieve the remaining entire text of a string where as ? can be used to retrieve only one character of a string.

Example 1: Lets assume that a machine name is "server1". If you query using ser*, then server1 is found. If you query using ser?, then the machine is not found. To find server1 using the ?, you must query using ser????, where each ? represents a character.

Example 2: Lets assume that the machine name is "CN=MACHINE1.OU=ENG.O=NOVELL.T=TREE". To find the machine, you can query by using "CN=MA*.OU=ENG.O=NOVELL.T=TREE" or CN=MA*. The machine name can be queried partly also. If you want to query by "O=novell.T=TREE", use * as "*O=novell.T=TREE".

The following table lists examples of wildcards usage:

| Example | Specifies to Include |
|---------------|--|
| * | All items |
| 164.99.* | All items starting with 164.99. |
| 164.9?.215.23 | All items starting with 164.9, followed by any character, and ending with ".215.23" |
| 164.96.215.23 | The single named item, in this case the inventoried server with the specified IP address |

4 Click *Run Selected Report*.

A status box appears displaying the progress of the report generation. When the report is generated, it appears in the viewer. Use the buttons on the toolbar to page through, print, or export the report.

NOTE: ZENworks Inventory report supports only the following double-byte character languages: German, English, Spanish, French, Portuguese, and Japanese. Other double-byte characters might not be displayed properly in the Inventory reports.

Printing an Inventory Report

To print a report:

- 1 **Generate and view the report.**
- 2 To change the default settings of the Printer, click the *Printer Setup* icon  and modify the settings.
- 3 Click the *Printer* icon .

Exporting an Inventory Report to a File

To export an inventory report to a file:

- 1** **Generate and view the report.**
- 2** On the toolbar, click the *Export Report* icon .
- 3** In the Export dialog box, specify the location and file format.

If you choose to export the Inventory report to a text file, in the Export to Text dialog box, select the *User defined* option and set the value to 16 because the data exported will be truncated if the value is less than 16.

If you want to export the Inventory report to an HTML file, you can select HTML 3.2 or HTML 4.0 (DHTML) file format. We recommend that you export to HTML 4.0 (DHTML) because the data exported to HTML 3.2 is not formatted properly.

If you want to export the Inventory report to a comma-separated value (.csv) file, do the following:

- 3a** Export the report to Microsoft* Excel.

NOTE: If you choose to export to .csv at this point, the report is not properly exported.

- 3b** Open the .xls file.
- 3c** Click *File*, then click *Save As*.
- 3d** In the *Save as type* field, choose *CSV (Comma delimited) (*.csv)*.
- 3e** Click *Save*.
- 4** Click *OK*.
- 5** Browse for and select the directory where you want to save the exported file.
- 6** Click *OK*.

Understanding User-Defined Reports

Using the Crystal Report Designer you can generate reports displaying information in the Inventory database.

Before generating the reports, you must ensure that the report file (.rpt) is created using Crystal Report Designer 8.0/8.5. For more information on how to create a .rpt file, see the Crystal Report documentation.

IMPORTANT: Except for the Software Inventory reports, you can use any Inventory report as a template to create a report.

To generate the User-defined Inventory report:

- 1** On the machine where you are designing the report, set the ODBC DSN name to ZenInventory.

To set the ODBC name:

- 1a** Click *Start*, click *Settings*, then click *Control Panel*.
- 1b** Double-click *ODBC Data Sources (32 Bit)*, then click *Add*.
- 1c** Select the ODBC driver for the database you want to connect to.
- 1d** Click *Finish*.

- 1e Specify the Data Source name as ZenInventory and specify the details.

NOTE: If you want to specify a data source name other than ZenInventory, you must configure the ODBC name on each of the machines where you invoke user-defined reports through ConsoleOne.

- 2 After you have designed the report, place the report in the
\`consoleone\version\reporting\canned\novellreporting\zeninventory\locale` directory.

Where *locale* can be EN for English language reports, FR for French language reports, PT_BR for Portuguese-Brazilian language reports, DE for German language reports, and ES for Spanish language reports. The non-English reports are displayed based on the respective locale of the machine.

- 3 Set the values in the `userreports.ini` file in the `consoleone\version\bin` directory. The `userreports.ini` file must contain the following values:

```
#[ReportName] <actual name of the report file without the .rpt extension>
#DisplayName=User Defined Report's display name
#Param1=Constant,Display name,<if combo then {val-1|val-2|val-3}>
#<where Param1 is the internal name of the parameter as stored in the .rpt file>
#<Constants are 1, 2 and 3 for Combo selection, text field and numeric field respectively>
```

For example, you can set the value as given below:

```
[ListSystemInformation]DisplayName=System Information
Role=1,Role,{2|3|5}
IPAddress=2,IP Address
DNName=2,Distinguished Name
DNTree=2,Distinguished Tree
DNSName=2,DNS Name
[ListMemory]
DisplayName=Memory
Role=1,Role,{2|3|5}
IPAddress=2,IP Address
DNName=2,Distinguished Name
DNTree=2,Distinguished Tree
DNSName=2,DNS Name
MemoryLowerLimit=3,Memory Lower Bound
```

- 4 After you set the values in the `userreports.ini` file, the User Defined Report is displayed in the Inventory Reports tree. You can specify multiple reports in the `userreports.ini` files.

NOTE: If the `userreports.ini` file is empty, the user cannot view the User Defined Reports in the Inventory Reports tree.

- 5 Click *Run Selected Report*.

17.1.5 Quickly and Easily Viewing the Inventory Data Using Quick Reports

In ZENworks 7 Workstation Inventory, provides a new tool called Quick Reports to easily retrieve and view the data from the ZENworks Inventory database. Each Quick Report contains a list of inventory attributes and a query that you define using the Quick Report wizard.

The following sections provide more information about working with Quick Report:

- ◆ “Invoking the Quick Report Wizard” on page 664
- ◆ “Creating a Quick Report” on page 664
- ◆ “Modifying an Existing Quick Report” on page 668
- ◆ “Viewing the Data Retrieved by the Quick Report” on page 669
- ◆ “Deleting a Quick Report” on page 671
- ◆ “Configuring the Inventory Database” on page 673
- ◆ “Working with the Query Results Window” on page 673

Invoking the Quick Report Wizard

Invoke the Quick Report Wizard using any of the following methods:

- ◆ To invoke the Quick Report from a database object, right-click the database object, click *ZENworks Inventory*, then click *Quick Report*.
- ◆ To invoke the Quick Report from the ConsoleOne Tools menu, click *ZENworks Inventory*, then click *Quick Report*.

If you have already configured the Inventory database, the Quick Report wizard uses that database.

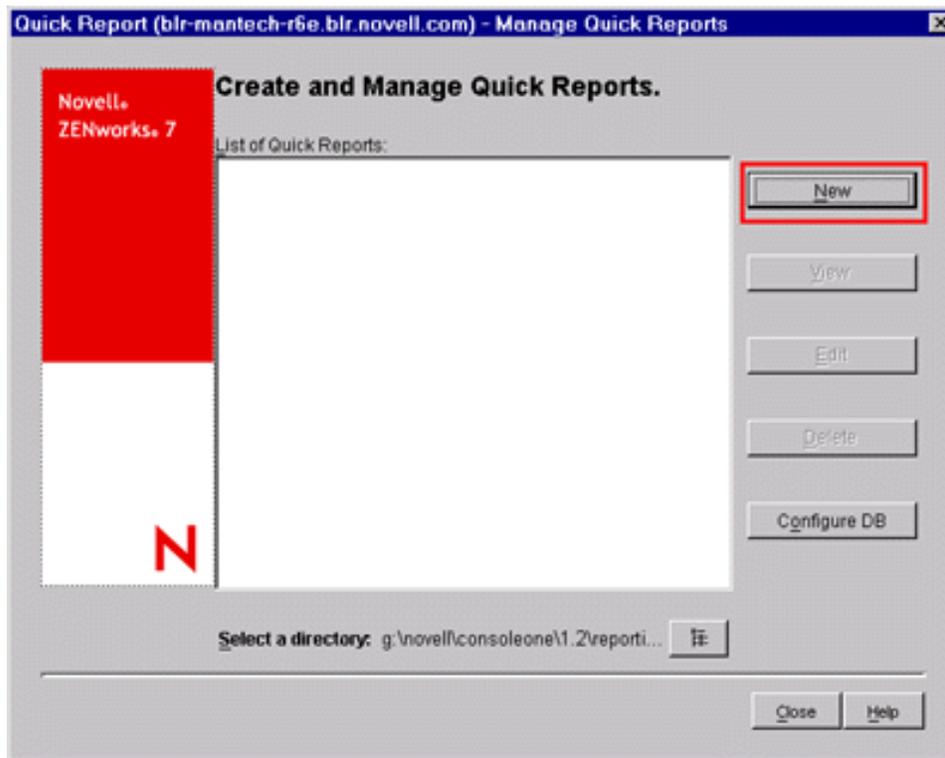
If you have not yet configured the Inventory database, the Quick Report wizard is displayed, and you can configure the database using the wizard. For more information, see “[Configuring the Inventory Database](#)” on page 673.

Creating a Quick Report

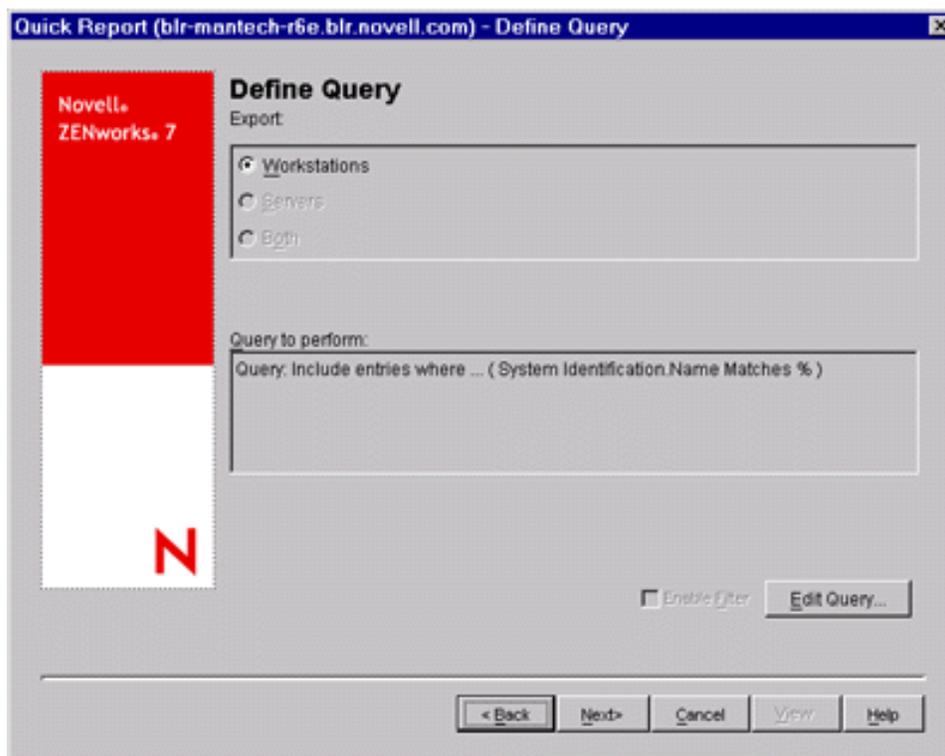
- 1 [Invoke the Quick Report wizard](#).

The Create and Manage Quick Reports page is displayed.

2 Click *New*.



3 In the Define Query page, define the query criteria and specify the scope for viewing the data from the Inventory database.



You can use either the default query or define a new query.

To use the default query, click *Next*. The Quick Report is created with the default query: System Identification.Name Matches %.

To define a new query:

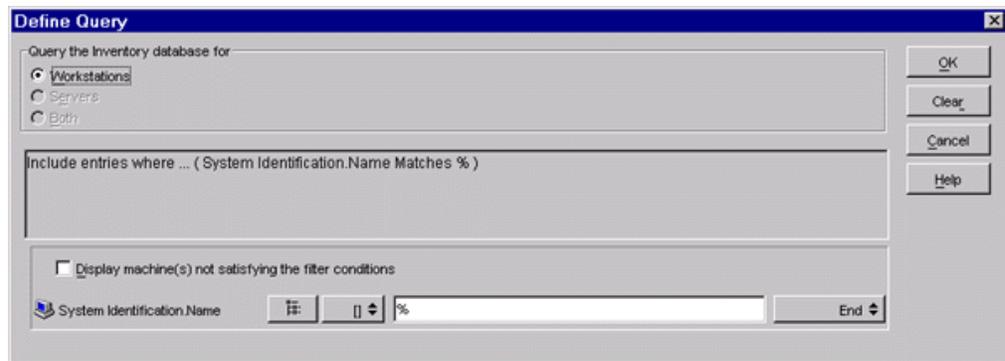
3a Select one of the following options.

- ♦ **Workstations:** Select this option to view the data satisfying the specified filter conditions for inventoried workstations. If you have only Desktop Management installed, this option is enabled by default and the other two options are unavailable.
- ♦ **Servers:** Select this option to view the data satisfying the specified filter conditions for inventoried servers. If you have only Server Management installed, this option is enabled by default and the other two options are unavailable.
- ♦ **Both:** Select this option to view the data satisfying the specified filter conditions for both inventoried servers and inventoried workstations. If you want to view data for inventoried workstations only, or for inventoried servers only, use one of the other query options. This option is available only if you have both ZENworks 7 Desktop Management and ZENworks 7 Server Management installed.

3b (Optional) If you want to apply the filter condition defined in the Define Query window, select the *Enable Filter* option.

This option is available only if you define the query using the following software classes and its attributes in the Define Query window: Software Group, Software Group File Information, Software Group Patch Information, Software, File Information, Patch Information, Exclude File Information, and Disk Usage.

3c Click *Edit Query* to change the query.

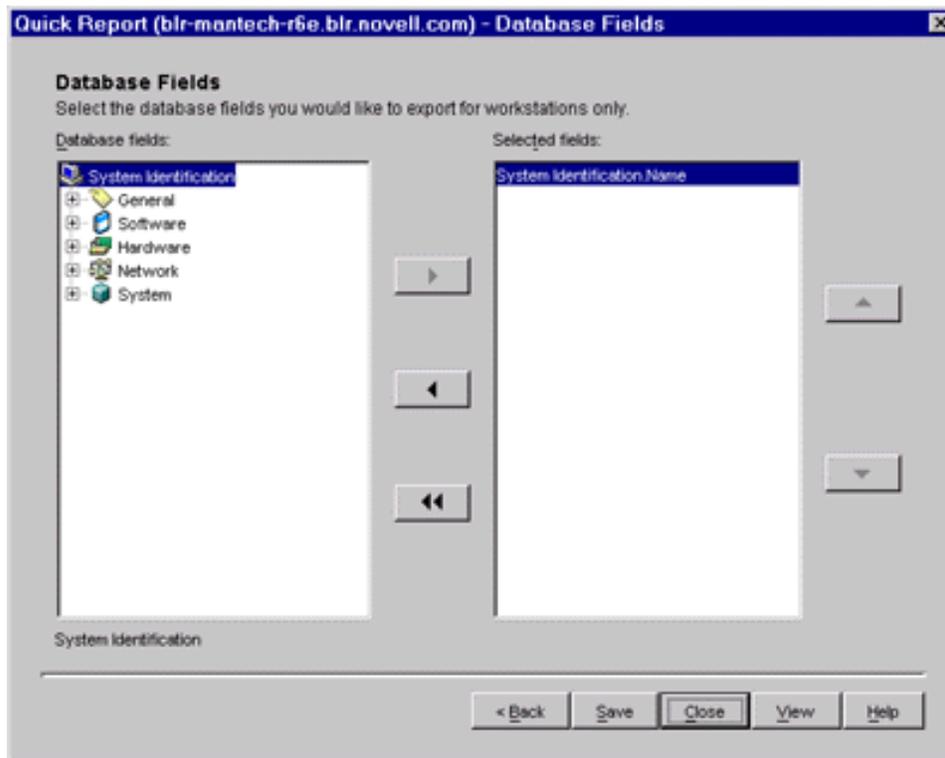


For detailed information on how to change the query, see [“Viewing Inventory Information of Inventoried Servers by Querying the Database”](#) on page 649.

3d Review the query and make changes as necessary. The Query to Perform pane displays the query you define.

3e Click *Next*.

4 In the Database Fields page, do the following:



4a From the *Database Fields* list, select the inventory attribute that you want to report.

By default, System Identification.Name is selected. You cannot deselect or change the order of this attribute.

4b Click  to add the selected inventory attribute to the *Selected Fields* list.

If you select a group attribute, all attributes of the group are added. For example, if you select the Software attribute, the Software attributes such as vendor name, product name, and version are included in the *Selected Fields* list.

4c To add an additional inventory attribute, repeat [Step 4a](#) and [Step 4b](#).

NOTE: You can change the order of the attributes using  and .

4d To view the report, click *View*.

The data is displayed in the Query Results window. For more information about the Query Results window, see [“Working with the Query Results Window”](#) on page 673.

4e To save the report, click *Save*.

IMPORTANT: Only the saved Quick Reports are listed on the Create and Manage Quick Reports page.

4f Click *Close*.

Modifying an Existing Quick Report

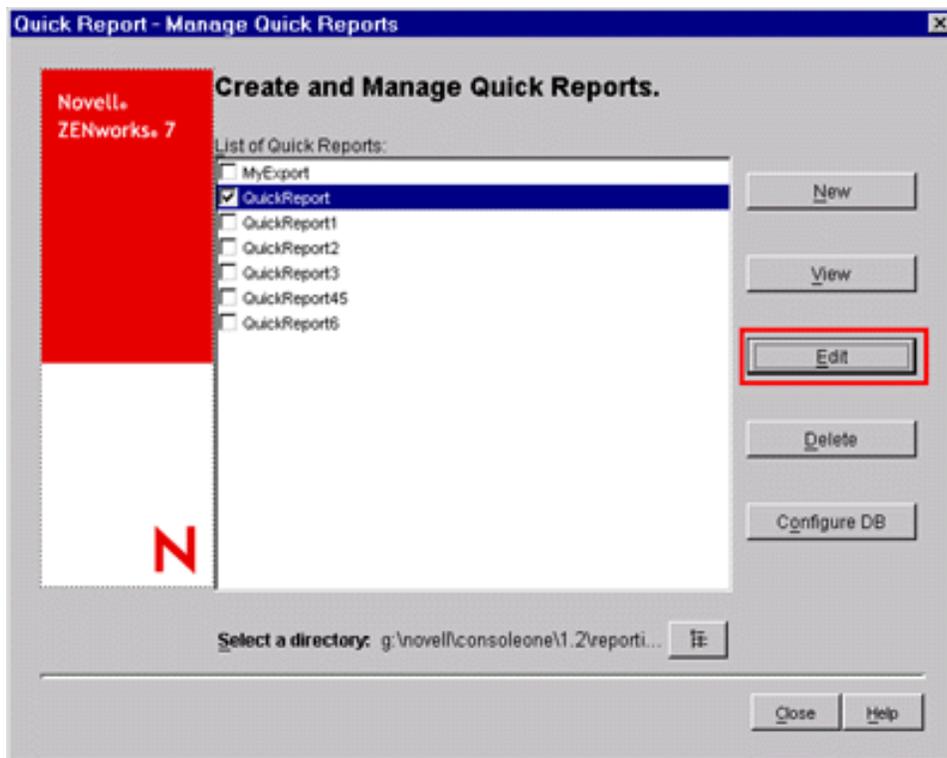
- 1 In the Create and Manage Quick Reports page, select the Quick Report that you want to modify from the list of Quick Reports.

You can modify only one Quick Report at a time.

By default, the list displays all the Quick Reports saved in the `ConsoleOne_installation_directory\consoleone\1.2\reporting\export` directory. To modify a Quick Report residing in another directory, click the *Browse* icon to browse and select the directory.

IMPORTANT: Only the saved Quick Reports are listed on the Create and Manage Quick Reports page.

- 2 Click *Edit*.



- 3 (Optional) In the Define Query page, change the existing query criteria and specify the scope for viewing the data from the Inventory database.
 - 3a Select one of the following options.
 - ♦ **Workstations:** Select this option to view the data satisfying the specified filter conditions for inventoried workstations.
 - ♦ **Servers:** Select this option to view the data satisfying the specified filter conditions for inventoried servers.
 - ♦ **Both:** Select this option to view the data satisfying the specified filter conditions for both inventoried servers and inventoried workstations.

3b (Optional) If you want to apply the filter condition defined in the Define Query window, select the *Enable Filter* option.

This option is available only if you define the query using the following software classes and its attributes in the Define Query window: Software Group, Software Group File Information, Software Group Patch Information, Software, File Information, Patch Information, Exclude File Information, and Disk Usage.

3c Click *Edit Query* to change the query.

For detailed information on how to change the criteria, see [“Viewing Inventory Information of Inventoried Servers by Querying the Database” on page 649](#).

3d Review the query and make changes as necessary. The Query to Perform pane displays the query you define.

3e Click *Next*.

4 (Optional) In the Database Fields page, do the following:

4a From the *Database Fields* list, select the inventory attribute that you want to report.

By default, System Identification.Name is selected. You cannot deselect or change the order of this attribute.

4b Click  to add the selected inventory attribute to the *Selected Fields* list.

If you select a group attribute, all attributes of the group are added. For example, if you select the Software attribute, the Software attributes such as vendor name, product name, and version are included in the Selected Fields list.

4c To add an additional inventory attribute, repeat [Step 4a](#) and [Step 4b](#).

NOTE: You can change the order of the attributes using  and .

4d To view the report, click *View*.

The report is displayed in the Query Results window. For more information about the Query Results window, see [“Working with the Query Results Window” on page 673](#).

4e To save the report, click *Save*.

IMPORTANT: Only the saved Quick Reports are listed on the Create and Manage Quick Reports page.

4f Click *Close*.

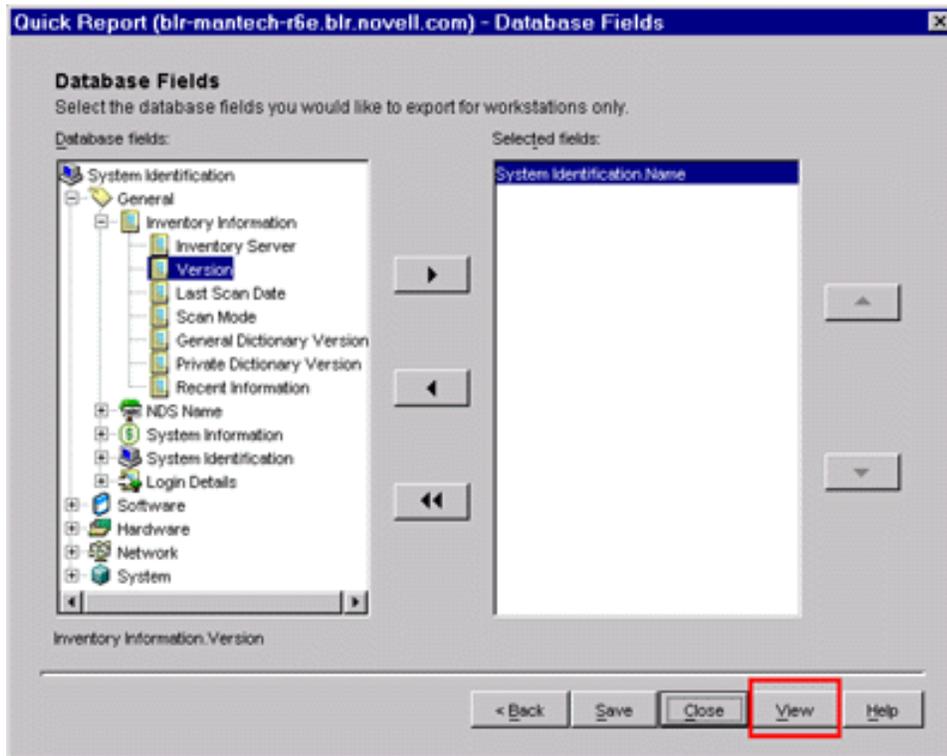
Viewing the Data Retrieved by the Quick Report

You can view the data retrieved by the Quick Report using any of the following methods:

- ◆ [“Viewing the Data While Creating or Modifying a Quick Report” on page 670](#)
- ◆ [“Viewing the Data of a Saved Quick Report” on page 670](#)

Viewing the Data While Creating or Modifying a Quick Report

- 1 In the Database Fields page, click *View*.



Viewing the Data of a Saved Quick Report

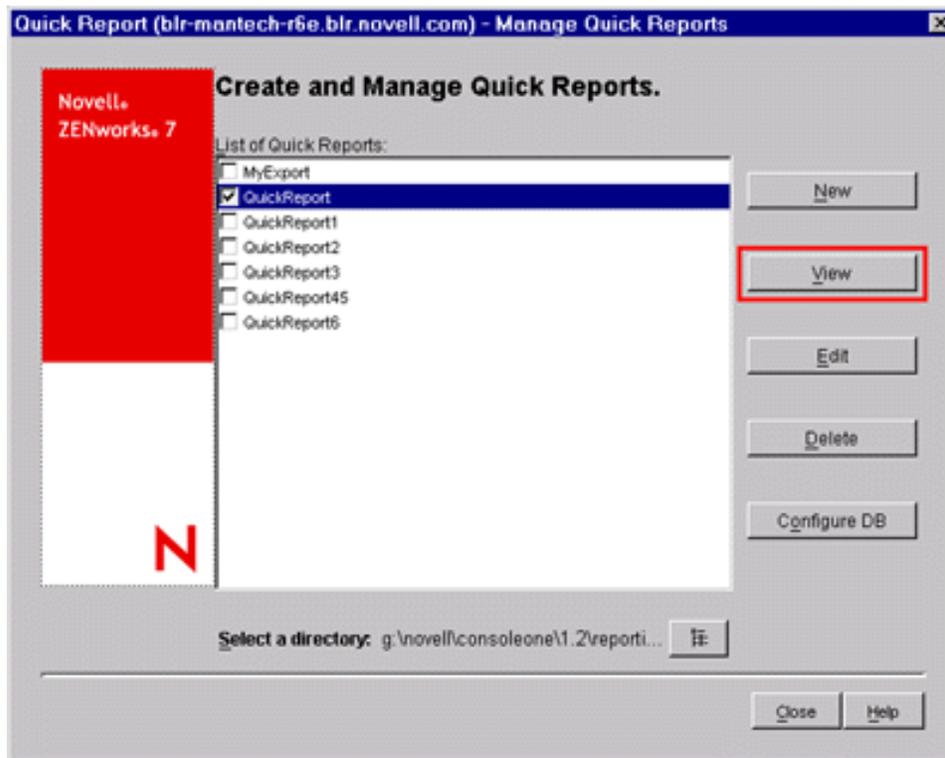
- 1 In the Create and Manage Quick Reports page, select the Quick Report that you want to view from the list of Quick Reports.

You can view only one Quick Report at a time.

By default, the list displays all the Quick Reports saved in the `ConsoleOne_installation_directory\consoleone\1.2\reporting\export` directory. To view a Quick Report residing in another directory, click the *Browse* icon to browse and select the directory.

IMPORTANT: Only the saved Quick Reports are listed on the Create and Manage Quick Reports page.

2 Click *View*.



The data is displayed in the Query Results window. For more information about the Query Results window, see [“Working with the Query Results Window”](#) on page 673.

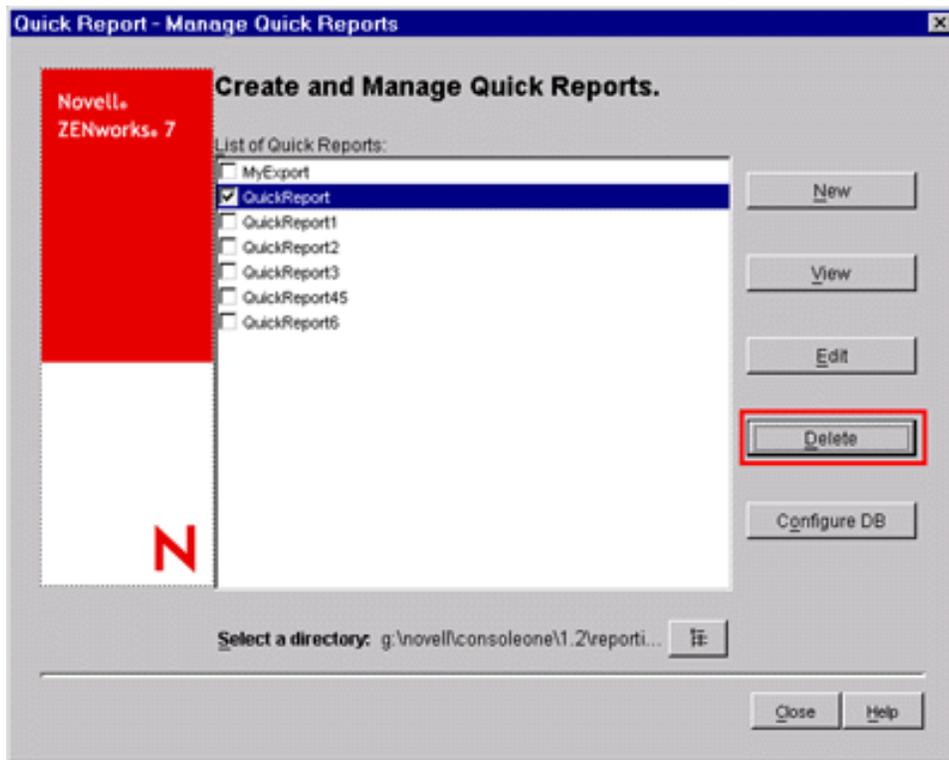
Deleting a Quick Report

1 In the Create and Manage Quick Reports page, select the Quick Report that you want to delete from the list of Quick Reports.

By default, the list displays all the Quick Reports saved in the `ConsoleOne_installation_directory\consoleone\1.2\reporting\export` directory. To delete a Quick Report residing in another directory, click the *Browse* icon to browse and select the directory.

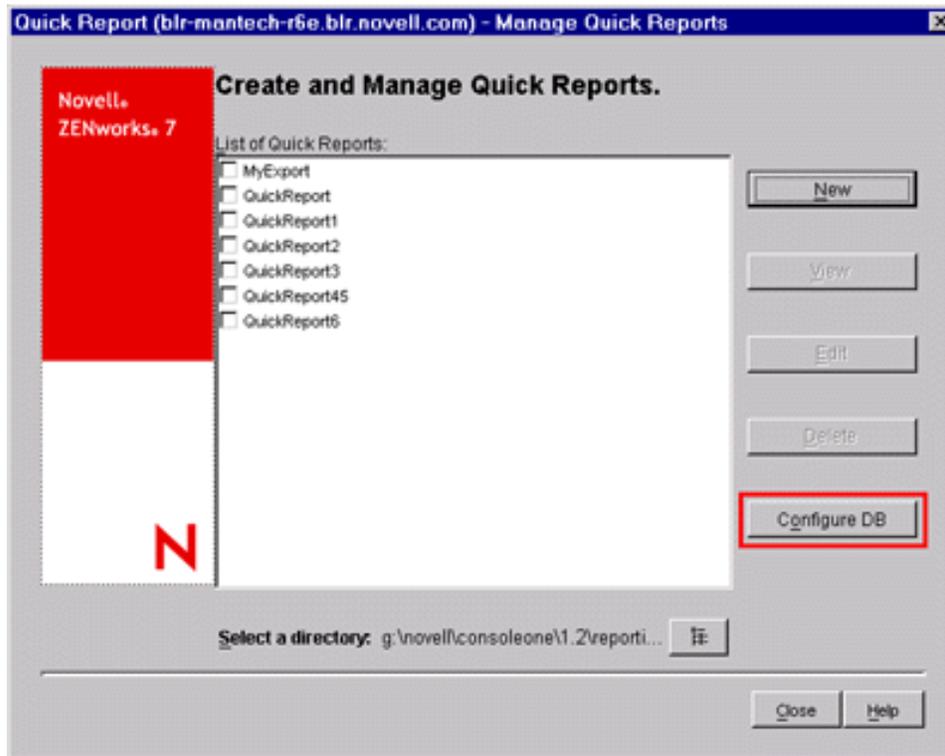
IMPORTANT: Only the saved Quick Reports are listed on the Create and Manage Quick Reports page.

2 Click *Delete*.



Configuring the Inventory Database

- 1 In Create and Manage Quick Reports page, click *Configure DB*.



The Configure ZENworks Database window is displayed.

- 2 Click *Browse* to select an existing ZENworks database object from the list.

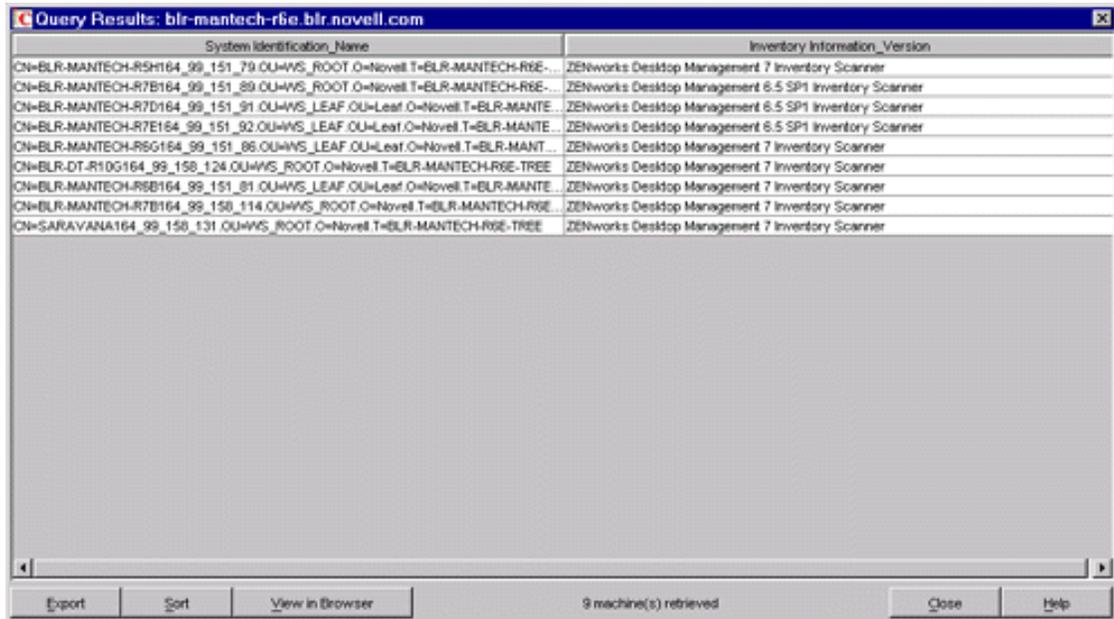
This database object contains the database settings such as the protocol, port in use by the database, and so forth.

- 3 Click *OK*.

Working with the Query Results Window

The Query Results window displays the data stored in the ZENworks Inventory database on querying the selected quick report. The Query Results window displays data for a maximum of 500 inventoried machines.

Figure 17-1 Query Results Window



The screenshot shows a window titled "Query Results: blr-mantech-r6e.blr.novell.com". It contains a table with two columns: "System identification_Name" and "Inventory information_Version". The table lists several entries with their respective system names and inventory scanner versions. At the bottom of the window, there are buttons for "Export", "Sort", "View in Browser", and "Close", along with a status bar indicating "9 machine(s) retrieved".

| System identification_Name | Inventory information_Version |
|---|---|
| CN=BLR-MANTECH-R5H164_99_151_79.OU=WYS_ROOT.O=Novell.T=BLR-MANTECH-R5E-TREE | ZENworks Desktop Management 7 Inventory Scanner |
| CN=BLR-MANTECH-R7B164_99_151_89.OU=WYS_ROOT.O=Novell.T=BLR-MANTECH-R5E-TREE | ZENworks Desktop Management 6.5 SP1 Inventory Scanner |
| CN=BLR-MANTECH-R7D164_99_151_91.OU=WYS_LEAF.OU=Leaf.O=Novell.T=BLR-MANTECH-R5E-TREE | ZENworks Desktop Management 6.5 SP1 Inventory Scanner |
| CN=BLR-MANTECH-R7E164_99_151_92.OU=WYS_LEAF.OU=Leaf.O=Novell.T=BLR-MANTECH-R5E-TREE | ZENworks Desktop Management 6.5 SP1 Inventory Scanner |
| CN=BLR-MANTECH-R5G164_99_151_96.OU=WYS_LEAF.OU=Leaf.O=Novell.T=BLR-MANTECH-R5E-TREE | ZENworks Desktop Management 7 Inventory Scanner |
| CN=BLR-DT-R10G164_99_158_124.OU=WYS_ROOT.O=Novell.T=BLR-MANTECH-R5E-TREE | ZENworks Desktop Management 7 Inventory Scanner |
| CN=BLR-MANTECH-R5B164_99_151_81.OU=WYS_LEAF.OU=Leaf.O=Novell.T=BLR-MANTECH-R5E-TREE | ZENworks Desktop Management 7 Inventory Scanner |
| CN=BLR-MANTECH-R7B164_99_158_114.OU=WYS_ROOT.O=Novell.T=BLR-MANTECH-R5E-TREE | ZENworks Desktop Management 7 Inventory Scanner |
| CN=SARAVANA164_99_158_131.OU=WYS_ROOT.O=Novell.T=BLR-MANTECH-R5E-TREE | ZENworks Desktop Management 7 Inventory Scanner |

You can perform the following operations in this window:

- ◆ Export entries to an xml or csv file.
- ◆ Sort the display of entries.
- ◆ View the data in a browser.

IMPORTANT: When you click *View in Browser*, the inventory data is exported in the XML format for rendering in the browser. Ensure that the browser, such as Microsoft Internet Explorer or Mozilla* Firefox, is the default application associated with the XML format.

If Internet Explorer is the default application associated with the XML format and it is already opened, and when you click *View in Browser*, you want the data to displayed in a new Internet Explorer window, do the following:

1. Invoke Windows Explorer on the machine running Quick Report.
 2. Click the *Tools* menu, then click *Folder Options*.
 3. In the Folder Options window, click the *File Types* tab.
 4. From the list of registered file types, select *XML*.
 5. In the *Details for 'XML' extension* pane, click *Advanced*.
 6. In the Edit File Type window, click *New*.
 7. Specify an action name, and in the *Application Used to Perform Action* field, type `Internet_Explorer_installation_directory\iexplore.exe -new %1`.
 8. Click *OK*.
 9. In Edit File Type window, select the newly created action from the *Actions* pane, and click *Set Default*.
 10. Click *OK*, then click *Close*.
- ◆ Stop the data retrieval process.

The Quick Report retrieves the data from the ZENworks Inventory database. You can stop the retrieval process by clicking *Stop* in the status bar of the Query Results dialog box.

The status bar displays the count of machines whose data has been retrieved. If you stop the process while the data for a single machine has not yet been completely retrieved, the Query Results dialog box displays the data retrieved until that time, but the status bar does not contain any message. And, if you stop the process while the data is being retrieved for multiple machines, the status bar displays the count of machines for which the data has been completely retrieved.

- ◆ Re-order the columns by dragging and dropping them.
- ◆ Re-size the columns.
- ◆ Select the entries by using the mouse or pressing Ctrl+A.
- ◆ Copy and paste the entries to the Clipboard by pressing Ctrl+C and Ctrl+V.

Exporting the Entries to an XML or a CSV File

- 1 Click *Export*.
- 2 In the Export Results dialog box, specify the XML or the CSV filename, and select the corresponding file type.

By default, the file is saved as `quick_report_name.xml` in the `ConsoleOne_installation_directory\consoleone\1.2\reporting\export` directory.

NOTE: If you specify a filename within double quotes, and without an extension or with an extension other than `.xml` or `.csv`, the file is stored in the comma-separated value (CSV) format irrespective of the file type you select.

Sorting the Display Ascending or Descending Order

- 1 Click *Sort*.
- 2 In the *Sort Items By* list, select the column by which you want to sort the entries.
- 3 Select either *Ascending* or *Descending*.
- 4 Configure the *Then By* drop-down lists.
- 5 Click *OK*.

For example, the Query Results window has the following entries:

| Product Name | Vendor Name |
|-----------------|-------------|
| Microsoft Word | Microsoft |
| Microsoft Excel | Microsoft |
| ZENworks | Novell |
| iPrint | Novell |
| GroupWise | Novell |
| Adobe Acrobat | Adobe |

If you want to sort the entries first by the vendor name in the ascending order and then sort all the Novell products in the ascending order, do the following:

- 1 Click the *Sort* button.
- 2 In the *Sort By* drop-down list, select *Vendor Name*.
- 3 Select the *Ascending* option.
- 4 In the *Then By* drop-down list, select *Product Name*.
- 5 Select the *Ascending* option.
- 6 Click *OK*.

The entries are displayed as shown below:

| Product Name | Vendor Name |
|-----------------|-------------|
| Adobe Acrobat | Adobe |
| Microsoft Excel | Microsoft |
| Microsoft Word | Microsoft |
| GroupWise | Novell |
| iPrint | Novell |
| ZENworks | Novell |

17.2 Exporting the Inventory Information

You can customize the inventory information you want to export from the ZENworks 7 Server Management Inventory database in to a comma-separated value (.csv) or an XML file.

You select the inventory components that should be exported, such as the Operating System Name and Version. You can further filter the inventoried servers whose attributes are exported. For example, you can export only those inventoried servers with a particular processor speed. The Data Export tool exports all inventoried servers satisfying these query conditions into a .csv or .xml file.

If you want to reuse the same data export settings for export, you can save the data export configurations.

The following sections help you use the Data Export tool:

- ◆ [Section 17.2.1, “Procedure to Export the Inventory Information,” on page 676](#)
- ◆ [Section 17.2.2, “Loading an Existing Configuration File,” on page 679](#)
- ◆ [Section 17.2.3, “Running the Data Export Program from the Inventory Server,” on page 680](#)
- ◆ [Section 17.2.4, “An Overview of XML and the Contents of an XML File,” on page 681](#)

17.2.1 Procedure to Export the Inventory Information

- 1 In ConsoleOne, select a container.

2 Invoke the Data Export tool.

- ♦ To invoke the Data Export tool from a database object, right-click the database object, click *ZENworks Inventory*, then click *Data Export*.
- ♦ To invoke the Data Export tool from the ConsoleOne Tools menu, you must first configure the Inventory database and then click *Tools*, click *ZENworks Inventory*, then click *Data Export*. For more information on how to configure the Inventory database, see [“Configuring the Inventory Database” on page 635](#).

3 Select *Create a New Database Query*.

This option lets you add a new query that defines the inventory components such as hardware, software, network, and others that you want to export. You can also specify the criteria to limit the inventoried servers and the database sites to be included in the query. Based on the inventory components and criteria you specify, the inventory information from the database is exported to a `.csv` or `.xml` file.

NOTE: If you want to load existing configuration settings for data export, select *Open a Saved Database Query*. This option lets you modify the settings for data export and then export the data to a `.csv` or `.xml` file. For more information, see [“Loading an Existing Configuration File” on page 679](#).

4 Click *Next*.

5 Specify the filter conditions for the inventoried servers.

5a Click *Edit Query*. For more information on how to define a query, see [“Viewing Inventory Information of Inventoried Servers by Querying the Database” on page 649](#).

5b (Optional) The *Enable Filter* option is available for selection only if you define the query using the software classes and its attributes of a supported category. Following are supported categories:

Category 1: Software Group, Software Group File Information, Software Group Patch Information, Software, File Information, and Patch Information

Category 2: Exclude File Information

Category 3: Disk Usage

The *Enable Filter* option is not available for selection if the query contains attributes belonging to different categories. For example, a query containing `software.name=word`, `softwaregroup.name=office`, and `diskusage.name=exe`.

If you want the results stored in `.csv` or `.xml` file to be filtered on the basis of the above query, select the *Enable Filter* check box.

5c Set the scope for exporting the information from the Inventory database.

If the ConsoleOne snap-ins and the Data Export tool have been installed for both ZENworks 7 Server Management and ZENworks 7 Desktop Management, the Data Export tool allows you to change the scope of exporting the inventory information.

By default, the *Servers* option is enabled. The query locates all inventoried servers satisfying the query expression. If ZENworks 7 Server Management and ZENworks 7 Desktop Management are installed in the same environment, the *Workstations*, the *Servers* and the *Both* options are available. When you select *Servers*, the query locates all inventoried servers satisfying the query expression. Choose *Both* to include all inventoried workstations and inventoried servers satisfying the query expression.

5d Review the query that you define.

5e Click *Next*.

6 Select the database fields from the list of database fields, then click *Add*.

If you select a group component, all subcomponents of the group are added. For example, if you select the Software component group, the subcomponents of Software such as vendor name, product name, and version are added.

7 Click *Next*.

8 View the data export settings.

8a Click *Save Configuration* to save the configurations settings to an `.exp` file. Specify the filename for the `.exp` file and then click *Save*.

The configuration file (`.exp`) contains the settings such as the inventory components you selected, and also the query formed for filtering the inventoried server data export. You create an `.exp` file so that you can reload the configuration settings and generate the `.csv` or `.xml` files any time you need to.

8b Click *Next*.

9 Select the machine from where you intend to perform the query.

9a Perform the Query from This Computer: Select *Perform the Query from This Computer* to run the data export processing from the computer. This option accesses the Inventory database on the specified database server and export the data in to a `.csv` or `.xml` file.

Perform the Query on a Remote Server: Select *Perform the Query on a Remote Server* to run the data export program from any server that has Server Inventory components installed.

Running the Data Export program from a server is recommended if you are exporting information from a large database with more than 10,000 inventoried servers or if you have specified complex queries with more than 20 database fields selected for exporting.

9b If you want to apply default encoding of the machine to the `.csv` or `.xml` file, select *Default Encoding*. The *Default Encoding* check box is selected by default. To apply Unicode encoding to the `.csv` or `.xml` file, select *Unicode Encoding*.

NOTE: If you create an `.exp` file to perform the data export from the local machine but use the same `.exp` to perform data export from a remote server and you want Unicode encoding, you must manually edit the `.exp` file and set the value of `DEExportEncode` to `UNICODE`.

9c Click *Next*.

10 Select an export option.

10a Select one the following options:

Export to CSV: Saves the inventory information in a `.csv` file.

Export to XML: Saves the inventory information in a `.xml` file.

10b Specify the path and the filename where you want to create the `.csv` or `.xml` file.

10c Click *Finish*.

If the configuration settings are not saved, you are prompted to save the changes

This generates the `.csv` or `.xml` file in the specified directory.

Open the `.csv` file in Microsoft Excel or any other CSV-supported viewer to view the exported data.

Open the `.xml` file in a XML viewer such as XML Spy. For more information, see [Section 17.2.4, “An Overview of XML and the Contents of an XML File,”](#) on page 681.

17.2.2 Loading an Existing Configuration File

You can load an existing configuration file (`.exp`). An `.exp` file contains the settings such as the inventory components you selected, and also the query formed for filtering the inventoried server data export.

After you load the `.exp` file, you can modify the settings for data export and then export the data to a `.csv` or `.xml` file.

To load existing configuration settings for data export:

- 1 Ensure that you have generated the data configuration files.

Complete the procedure outlined in [Section 17.2.1, “Procedure to Export the Inventory Information,”](#) on page 676. This procedure generates the `.csv` or `.xml` file and the data configuration files.

- 2 In ConsoleOne, select a container and invoke the Data Export tool using any of the following methods:

- ♦ To invoke the Data Export tool from a database object, right-click the database object, click *ZENworks Inventory*, then click *Data Export*.
- ♦ To invoke the Data Export tool from the ConsoleOne Tools menu, you must first configure the Inventory database and then click *Tools*, click *ZENworks Inventory*, then click *Data Export*. For more information on how to configure the Inventory database, see [“Configuring the Inventory Database”](#) on page 635.

- 3 Select *Open a Saved Database Query*, then click *Next*.

The default directory for `.exp` files is

`consoleone\consoleone_version\reporting\export`. Click *Browse* to open an existing `.exp` file.

If the `.exp` and `.cfg` files are invalid or are an older version, the data export will not proceed. The data export displays the number of servers and servers that satisfy the query and filter conditions for export.

- 4 Select a saved database query from the list of saved queries.

- 4a Select a saved database query from the list of saved queries. The list box displays the `.exp` files that are saved in `consoleone\consoleone_version\reporting\export`.

or

Click *Browse* to open an existing `.exp` file in any other location.

- 4b (Optional) If the `.exp` and `.cfg` files are invalid or are an older version, the data export will not proceed. The data export displays the number of servers and servers that satisfy the query and filter conditions for export.

If you want to modify the existing query, click *Edit* and modify the query and select the new database fields. For more information on how to define a query, see “[Viewing Inventory Information of Inventoried Servers by Querying the Database](#)” on page 649.

- 4c** Click *Next*.
- 5** To view the data export settings:
 - 5a** Click *Save Configuration* to save the configurations settings to an `.exp` file. Specify the filename for the `.exp` file and then click *Save*.

The configuration file (`.exp`) contains the settings such as the inventory components you selected, and also the query formed for filtering the inventoried server data export. You create an `.exp` file so that you can reload the configuration settings and generate the `.csv` or `.xml` files any time you need to.
 - 5b** Click *Next*.
- 6** Select the machine from where you intend to perform the query.
 - 6a** **Perform the Query from This Computer:** Select *Perform the Query from This Computer* to run the data export processing from the computer. This option accesses the Inventory database on the specified database server and export the data in to a `.csv` or `.xml` file.

Perform the Query on a Remote Server: Select *Perform the Query on a Remote Server* to run the data export program from any server that has Server Inventory components installed.

Running the Data Export program from a server is recommended if you are exporting information from a large database with more than 10,000 inventoried servers or if you have specified complex queries with more than 20 database fields selected for exporting.
 - 6b** If you want to apply default encoding of the machine to the `.csv` or `.xml` file, select *Default Encoding*. The *Default Encoding* check box is selected by default. To apply Unicode encoding to the `.csv` or `.xml` file, select *Unicode Encoding*.
 - 6c** Click *Next*.
- 7** Select an export option.
 - 7a** Select one the following options:
 - Export to CSV:** Saves the inventory information in a `.csv` file.
 - Export to XML:** Saves the inventory information in a `.xml` file.
 - 7b** Specify the path and the filename where you want to create the `.csv` or `.xml` file.
 - 7c** Click *Finish*.

17.2.3 Running the Data Export Program from the Inventory Server

Running the Data Export program from a server is recommended if you are exporting information from a large database with more than 10,000 inventoried servers or if you have specified complex queries with more than 20 database fields selected for exporting.

To run the data export program from the server:

- 1** Ensure that you have generated the data configuration files.

Follow the Step 1 to Step 5 as outlined in [Section 17.2.1, “Procedure to Export the Inventory Information,”](#) on page 676 and ensure that you save the settings in the `.exp` file.

When you save an `.exp` file, a corresponding data configuration file is created in the same directory with the same filename as the `.exp` file and with the `.cfg` file extension.

- 2 Click *Perform the Query on a Remote Server* to run the data export program from any server that has Server Inventory components installed, then click *Finish*.
- 3 Copy the `.exp` file and `.cfg` file to the server.

These two files should exist in the same directory on the Inventory server.

From the server console, run `dbexport.ncf` on NetWare servers, `dbexport.bat` on Windows servers. To do so, enter:

```
DBEXPORT "configuration_filename.exp" "csv_filename.csv"
```

where `configuration_filename.exp` is an existing file that contains the data export settings. You must enter the `configuration_filename.exp` and the `csv_filename.csv` filenames within double quotes. The data exported from the database is stored in `csv_filename.csv`.

The corresponding `.cfg` file for the `.exp` file should be in the same folder as the `.exp` file. The `.cfg` file contains the list of the database attributes to be exported.

If the `.exp` and `.cfg` files are invalid or are older versions, the data export will not proceed. The data export displays the number of inventoried servers that satisfy the query and filter conditions for export.

17.2.4 An Overview of XML and the Contents of an XML File

Server Inventory allows you to export the inventory information from the Inventory database into an Extensible Markup Language (`.xml`) file by using the Data Export tool.

XML is a markup language that provides a format for describing structured data. An XML document is a text-based format. The XML source is made up of XML elements. The XML tags are not predefined and you must define your own tags.

For more information about XML, see the [World Wide Web Consortium \(W3C\) XML Activity and Information web site \(http://www.w3.org/XML\)](http://www.w3.org/XML).

A sample `.xml` file is as follows:

```
<?xml version="1.0" encoding='UTF-8' ?>
<!DOCTYPE InventoryInformation [<!ELEMENT Attribute (value)>
<!ATTLIST Attribute
    name CDATA #REQUIRED
    type (custom | regular) #REQUIRED
    units CDATA #IMPLIED
>
<!ELEMENT Class (Attribute*)>
<!ATTLIST Class
    name CDATA #REQUIRED
    instance CDATA #REQUIRED
>
<!ELEMENT InventoryInformation (Machine+)>
<!ELEMENT Machine (Class+)>
```

```

<!ATTLIST Machine
    name CDATA #REQUIRED
>
<!ELEMENT value (#PCDATA)>]
>
<InventoryInformation>
  <Machine name="blr-stl-zen1.blr.novell.com">
    <Class name="Processor" instance="1">
      <Attribute name="Current Clock Speed" type="regular" units="MHz">
        <value>2800</value>
      </Attribute>
      <Attribute name="Processor Family" type="regular">
        <value>"Intel (R) Xeon (TM)"</value>
      </Attribute>
    </Class>
    <Class name="IP" instance="1">
      <Attribute name="IP Address" type="regular">
        <value>164.99.163.9</value>
      </Attribute>
      <Attribute name="Subnet Mask" type="regular">
        <value>255.255.252.0</value>
      </Attribute>
    </Class>
  </Machine>
</InventoryInformation>

```

XML uses a Document Type Definition (DTD) to describe the data. The DTD is embedded within the XML document.

A DTD lists the elements, attributes, and entities contained in a document and also, defines the relationship between the elements and attributes.

Following is the DTD embedded in the preceding sample xml file:

```

<?xml version="1.0" encoding='UTF-8' ?>
<!DOCTYPE InventoryInformation [<!ELEMENT Attribute (value)>
<!ATTLIST Attribute
    name CDATA #REQUIRED
    type (custom | regular) #REQUIRED
    units CDATA #IMPLIED
>
<!ELEMENT Class (Attribute*)>
<!ATTLIST Class
    name CDATA #REQUIRED
    instance CDATA #REQUIRED
>
<!ELEMENT InventoryInformation (Machine+)>
<!ELEMENT Machine (Class+)>
<!ATTLIST Machine

```

```

    name CDATA #REQUIRED
  >
<!ELEMENT value (#PCDATA)>]
>

```

Table 17-3 explains the elements used in the sample XML file:

Table 17-3 Sample XML File Elements

| Elements Used in the Sample XML File | Description |
|--------------------------------------|-----------------------------|
| Class | Device name |
| Type | Custom or Regular attribute |
| Units | Unit information |
| Instance | Device instance count |

17.3 Retrieving Inventory information from the Inventory Database Without Using the CIM Schema

ZENworks 7 Server Management provides easy-to-use Inventory database views that allow you to retrieve inventory information from the Inventory database without using the CIM schema.

The Inventory views are predefined device-specific views that are automatically created in the Inventory database after you install the Server Inventory component of ZENworks 7 Server Management.

The nomenclature for the Inventory views is *database_schema_name.zen_devicename*. For example, *mw_dba.zen_processor*.

Inventory views that are associated with enums have localized views. For example, *mw_dba.zen_processor_ja* is the Japanese view for the Processor.

The following sections provide information about the various Inventory views and how to use them:

- ♦ [Section 17.3.1, “List of Inventory Views,” on page 683](#)
- ♦ [Section 17.3.2, “How to Use the Inventory Views,” on page 709](#)

17.3.1 List of Inventory Views

Table 17-4 ZENworks Inventory Views and their functionality

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|-------------------------|---------------------------------|------------|------------------------------------|------------------------|
| <i>mw_dba.zen_asset</i> | Retrieves the Asset information | SystemName | DNS name of the inventoried server | Yes |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|---------------------|-----------------------------------|----------------------|---|------------------------|
| | | Manufacturer | Name of the manufacturer | |
| | | Model | Model of the computer system | |
| | | SerialNumber | Serial number of the computer system assigned by manufacturer | |
| | | Tag | Unique identifier of system information | |
| | | ManagementTechnology | Technology available on the inventoried server such as DMI, WMI, and others | |
| | | AssetTag | Asset tag number that the ROM-based setup program creates | |
| | | ModelNumber | Model number of the computer system | |
| mw_dba.zen_battery | Retrieves the Battery information | Name | Device name for the battery, for example, Duracell* DR-36 | Yes |
| | | Chemistry | The battery chemistry, for example, lithium-ion or nickel metal hydride | |
| | | DesignCapacity | The design capacity of the battery in mWatt-hours | |
| | | DesignVoltage | The design voltage of the battery in mVolts | |
| | | SmartBatteryVersion | The Smart Battery Data Specification version number supported by this battery | |
| | | InstallDate | The battery manufacture date | |
| | | Manufacturer | The name of the company that manufactured the battery | |
| | | SerialNumber | The serial number for the battery | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|------------------------|--|------------------|--|------------------------|
| mw_dba.zen_bios | Retrieves the BIOS information | Caption | BIOS label | Yes |
| | | InstallDate | The manufacturing date of the BIOS | |
| | | SerialNumber | Serial number of the computer, assigned during manufacture | |
| | | Version | Version or revision level of the BIOS | |
| | | Manufacturer | BIOS vendor name | |
| | | PrimaryBIOS | True state indicates Primary BIOS | |
| | | BIOSIDBytes | Byte in the BIOS that indicates the computer model | |
| | | Size | Size of the BIOS | |
| mw_dba.zen_bus | Retrieves the Bus information | BusType | Bus type indicates PCI, ISA, and others | Yes |
| | | BusName | Bus name | |
| | | BusDescription | Bus description | |
| | | BusVersion | Version of the bus supported by the motherboard | |
| | | DeviceID | The unique hexadecimal ID for the specific bus | |
| mw_dba.zen_cachememory | Retrieves the Cache memory information | ErrorMethodology | Error correction scheme supported by this cache component, for example, Parity/ Single Bit ECC/ MultiBit ECC | Yes |
| | | Level | Indicates the cache level; internal cache that is built in to the microprocessors; external cache that is between the CPU and DRAM | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|---------------------|-----------------------------------|--------------------|--|------------------------|
| | | WritePolicy | Indicates the two different ways (Write-Back and Write-Through Cache) that the cache can handle to write to the memory | |
| | | CacheType | Defines the system cache type, for example, Instruction, Data, Unified | |
| | | LineSize | Size in bytes of a single cache bucket or line | |
| | | ReplacementPolicy | Algorithm that the cache uses to determine which cache lines or buckets should be reused | |
| | | ReadPolicy | Indicates whether the data cache is for read operations | |
| | | Associativity | Defines the system cache associativity (directmapped, 2-way, 4-way) | |
| | | Speed | Speed of this System Cache module in nanoseconds | |
| | | Capacity | Size of the data store where the cache information is kept | |
| mw_dba.zen_cdrom | Retrieves the CDROM information | DeviceID | Drive letter allocated for the CD on the inventoried server | No |
| | | Manufacturer | Vendor name of the CD | |
| | | Description | Description of the CD | |
| | | Caption | Caption of the CD | |
| mw_dba.zen_chassis | Retrieves the Chassis information | AssetTag | Asset tag number of the system chassis. For example, S11127 | Yes |
| | | NumberOfPowerCords | Total number of power cords attached to a system chassis | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|---------------------------------|---|---------------------|---|------------------------|
| | | ChassisType | Represents whether the system chassis is a laptop, desktop, notebook, docking station and so on | |
| | | Manufacturer | Name of the system chassis manufacturer. For example, Compaq | |
| | | SerialNumber | Manufacturer's number used to identify a system chassis. For example, 53R661S | |
| | | Tag | Unique ID of the system chassis attached to a particular inventoried server. For example, System Enclosure 0 | |
| | | Version | Version number of the system chassis | |
| mw_dba.zen_computerinformation | Retrieves the computer information | ComputerName | Name of the inventoried server as represented in eDirectory, such as the fully qualified DN of the inventoried server | No |
| | | PrimaryOwner | The name of the primary user or owner of this system | |
| | | PrimaryOwnerContact | The phone number of the primary user of this system | |
| mw_dba.zen_currentlogin_details | Retrieves the current login information | CurrentUser | User logged in to the Primary eDirectory tree when the inventoried server was scanned | No |
| mw_dba.zen_disk | Retrieves the disk information | RemovableDisk | Removable disk | Yes |
| | | Manufacturer | Vendor name of the disk | |
| | | Description | Description of the disk | |
| | | PhysicalCylinders | Number of cylinders | |
| | | PhysicalHeads | Number of heads | |
| | | SectorsPerTrack | Removable disk drive sectors per track | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|---------------------------|---|-----------------------------|---|------------------------|
| | | Capacity | Capacity of the removable disk | |
| mw_dba.zen_diskusage | Retrieves the disk usage information | FileExtension | The file extension for which the disk usage is scanned for. | No |
| | | TotalDiskUsage | Total disk usage for all the files of the specified extension. | |
| mw_dba.zen_displayadapter | Retrieves the display adapter information | Description | Description of the display adapter | Yes |
| | | VideoMemoryType | The type of video memory for this adapter, for example, VRAM/SRAM/DRAM/EDO RAM | |
| | | MaxMemorySupported | Maximum memory that the display adapter supports for VIDEO RAM | |
| | | CurrentBitsPerPixel | Number of adjacent color bits for each pixel | |
| | | CurrentHorizontalResolution | Number of horizontal pixels shown by the display | |
| | | CurrentVerticalResolution | Number of vertical pixels shown by the display | |
| | | MaxRefreshRate | Maximum refresh rate of the monitor for redrawing the display, measured in Hertz | |
| | | MinRefreshRate | Minimum refresh rate of the monitor for redrawing the display, measured in Hertz | |
| | | VideoArchitecture | The architecture of the video subsystem in this system, for example, CGA/VGA/SVGA/8514A | |
| | | NumberOfColorPlanes | Number of color planes supported by the video system | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|------------------------------|--|-------------------|---|------------------------|
| | | ChipSet | Chip set used by the controller to compare system capabilities | |
| | | DACType | Digital-to-Analog converter type | |
| | | ProviderName | Vendor name | |
| mw_dba.zen_displaydriver | Retrieves the display driver information | Manufacturer | Name of the display driver manufacturer | Yes |
| | | Version | Version number of the display driver | |
| | | InstallDate | Install date of the display driver | |
| | | IsShadowed | If True, the display driver is currently being shadowed | |
| mw_dba.zen_distinguishedname | Retrieves the distinguished name | DistinguishedName | Distinguished name | No |
| | | Tree | eDirectory tree name | |
| mw_dba.zen_dma | Retrieves the DMA information | Description | Name of the logical device that is currently using this DMA channel | Yes |
| | | DMACHannel | Number of the Direct Memory Access (DMA) channel that a computer uses for transferring data to and from devices quicker than from computers without a DMA channel | |
| | | Availability | Indicates whether Virtual Direct Memory Access (DMA) is supported | |
| | | BurstMode | A data transmission mode in which data is sent faster than normal | |
| mw_dba.zen_dnsname | Retrieves the DNS name | HostName | DNS name of the inventoried server | No |
| mw_dba.zen_floppy | Retrieves the floppy information | DeviceID | The floppy name representing the floppy | No |
| | | Manufacturer | Vendor name | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|-----------------------------|---|--------------------------|---|------------------------|
| | | Description | Floppy drive description | |
| | | PhysicalCylinders | Total number of cylinders or tracks on the floppy. | |
| | | PhysicalHeads | Floppy drive R/W heads | |
| | | SectorsPerTrack | Floppy drive sectors per track | |
| | | Capacity | Floppy drive capacity | |
| mw_dba.zen_inventoryscanner | Retrieves the inventory scanner information | LastScanDate | The date when the Scanner was last scanned. Stored as milliseconds so it can be read and displayed in any appropriate date format | Yes |
| | | InventoryServer | Name of the Inventory server to which the scans are sent. It is not the complete DN of the server name | |
| | | Version | Version of the Scanner running on the inventoried server | |
| | | ScanMode | The management technology used by the Scanner, such as WMI or DMI, for scanning the computer system | |
| | | RecentInformation | Latest inventory information | |
| | | generaldictionaryversion | Version of the General dictionary | |
| | | privatedictionaryversion | Version of the Private dictionary | |
| mw_dba.zen_ipaddress | Retrieves the IP address | Address | The unique address assigned to a computer on an IP Internet | No |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|-----------------------------|------------------------------------|----------------------|--|------------------------|
| | | SubnetMask | The subnet mask of the inventoried server paired with an IP address specifies to an IP router which octets or bits in the IP address are the network ID and which octets or bits are the node ID | |
| mw_dba.zen_ipxaddress | Retrieves the IPX address | Address | The IPX address of the inventoried server | No |
| mw_dba.zen_irq | Retrieves the IRQ information | IRQNumber | Number of the Interrupt Request Line (IRQ), from 0 to 15 | Yes |
| | | Availability | Availability of the specific IRQ channel | |
| | | TriggerType | IRQ Trigger type | |
| | | Shareable | If True, the system IRQ can be shared across devices | |
| mw_dba.zen_keyboard | Retrieves the keyboard information | KeyboardLayout | Layout of the keyboard | No |
| | | KeyboardSubtype | Type of the keyboard | |
| | | KeyboardDescription | Description of the keyboard, such as IBM Enhanced 101 or 102 keys | |
| | | NumberOfFunctionKeys | Total number of function keys | |
| | | KeyboardDelay | Delay before the repeat of a key | |
| | | TypematicRate | Rate of processing the keys | |
| mw_dba.zen_lastlogindetails | Retrieves the last login details | LastUser | User most recently logged in to the Primary eDirectory tree through Novell Client when the inventoried server was scanned | No |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|--------------------------------|--|--------------------|--|-------------------------------|
| mw_dba.zen_macaddresses | Retrieves the MAC address | MACAddress | Unique node address permanently coded in the network adapter that identifies a specific computer on a network | No |
| mw_dba.zen_memory | Retrieves the memory information | TotalMemory | Total memory of the inventoried server | No |
| mw_dba.zen_microsoftdomainname | Retrieves the Microsoft domain name | DomainName | Domain name of the inventoried server | No |
| mw_dba.zen_internalmodem | Retrieves the internal modem information | Name | Identifying information of the modem | No |
| | | Description | Additional information about the modem | |
| | | ProviderName | Name of the vendor | |
| | | DeviceID | Special hexadecimal string identifying the modem type | |
| mw_dba.zen_monitor | Retrieves the monitor information | DeviceID | Unique ID of a desktop monitor that is attached to an inventoried server. For example, DesktopMonitor1. | No |
| | | ModelID | Unique ID of a model of the monitor. It is a combination of the Manufacturer ID and Product ID. For example, DELA001. | |
| | | MonitorDescription | Description of the monitor. | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|------------------------|---------------------------------------|---------------------|--|------------------------|
| | | NominalSize | <p>A number representing the diagonal width of the monitor (the distance from one corner of the screen to the opposite corner of the screen)</p> <p>For example, 17"</p> <p>You can customize the scan of the nominal size of the monitor by configuring the HWRules ini file using the Server Inventory policy.</p> | |
| | | ViewableSize | <p>A number representing the diagonal width of the screen image excluding the black borders around the image's edge</p> <p>For example, 15.8"</p> | |
| | | ManufacturedDate | Year in which the monitor was manufactured | |
| | | MonitorSerialNumber | <p>Manufacturer's number used to identify a monitor</p> <p>For example, 23DDC24N9067</p> | |
| | | Manufacturer | <p>Name of the monitor's manufacturer</p> <p>For example, DELL Computer Corp</p> | |
| | | Model | <p>Product name of the monitor given by the manufacturer</p> <p>For example, DELL E771a</p> | |
| mw_dba.zen_motherboard | Retrieves the motherboard information | Description | General description of the motherboard | No |
| | | Manufacturer | Name of the motherboard manufacturer | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|---------------------------|---|------------------|--|------------------------|
| | | Version | Version of the motherboard | |
| | | NumberOfSlots | The number of expansion slots in the motherboard for adding more memory, graphic capabilities, and support for special devices | |
| mw_dba.zen_mouse | Retrieves the mouse information | MouseType | Mouse type | Yes |
| | | MouseName | Identifying information of the mouse | |
| | | NumberOfButtons | Number of buttons on the mouse | |
| | | IRQNumber | Interrupt assigned to this device | |
| mw_dba.zen_mousedriver | Retrieves the mouse driver information | DriverName | Name of the mouse driver | No |
| | | DriverVersion | Version number of the mouse driver | |
| mw_dba.zen_NetworkAdapter | Retrieves the network adapter information | Caption | Network adapter caption | Yes |
| | | Description | Network adapter description | |
| | | InstallDate | Install date of the network adapter | |
| | | Name | Network adapter name | |
| | | PermanentAddress | Node address stored permanently in the adapter | |
| | | MACAddress | The MAC address stored in the network adapter | |
| | | MaxSpeed | Rate at which the data is transferred over the LAN | |
| | | AdapterType | Type of network adapter, such as FDDI or token ring | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|---------------------------------|--|------------------|---|------------------------|
| | | ProviderName | Name of the provider | |
| mw_dba.zen_NetworkAdapterDriver | Retrieves the network adapter driver information | Description | Description of the network adapter driver installed on the inventoried server For example, IBM 10/100 Ethernet adapter, EN-2420Px Ethernet adapter | No |
| | | Name | Name of the network adapter driver | |
| | | Version | Version of the network adapter | |
| mw_dba.zen_parallelport | Retrieves the parallel port information | PortName | The logical name of the input-output device on this parallel port, under this operating environment | Yes |
| | | HasDMASupport | If True, DMA is supported | |
| | | PortAddress | Base I/O address for this parallel port | |
| | | IRQNumber | IRQ number of the parallel port | |
| mw_dba.zen_power supply | Retrieves the power supply information | Description | Expanded description of the input voltage capability for this power supply | No |
| | | TotalOutputPower | Attribute value that represents the total output power of the power supply | |
| mw_dba.zen_processor | Retrieves the processor information | DeviceID | Special hexadecimal string identifying the processor type | Yes |
| | | Description | Additional information about the processor | |
| | | Role | Type of processor such as central processor, math coprocessor, and others | |
| | | Family | Identification of the processor family such as Pentium II, Pentium III, and others | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|-------------------------|---|------------------------|--|------------------------|
| | | OtherFamilyDescription | Additional description about the Processor Family, such as Pentium Processor with MMX technology | |
| | | UpgradeMethod | The method by which this processor can be upgraded, if upgrades are supported | |
| | | MaxClockSpeed | Maximum clock speed of the processor | |
| | | CurrentClockSpeed | Current clock speed of the processor | |
| | | Stepping | Single-byte code characteristic provided by microprocessor vendors to identify the processor model | |
| mw_dba.zen_serialport | Retrieves the serial port information | PortName | The logical name of the I/O device on this serial port, under this operating environment | No |
| | | PortAddress | Base input-output address for this serial port | |
| | | IRQNumber | IRQ number of the serial port | |
| mw_dba.zen_soundadapter | Retrieves the sound adapter information | Name | Label of the multimedia card | No |
| | | Description | Description of the multimedia component for the server | |
| | | ProviderName | Name of the provider | |
| mw_dba.zen_systemslot | Retrieves the system slot information | SlotDescription | Card currently occupying this slot | No |
| | | MaxDataWidth | Maximum bus width of cards accepted in the slot | |
| | | ThermalRating | Maximum thermal dissipation of the slot in milliwatts | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|----------------------|--|----------------------|---|------------------------|
| mw_dba.zen_unixOS | Retrieves the UNIX operating system information | Type | Operating system of the inventoried server | Yes |
| | | Caption | Operating system name | |
| | | CodePage | Language code page of the operating system | |
| | | Version | Version number of the operating system | |
| | | InstallDate | Install date of the operating system | |
| | | VirtualMemorySize | Total number of bytes in the virtual address space of the calling process | |
| | | VisibleMemorySize | Total memory as reported by the operating system | |
| | | ProviderName | Name of the provider | |
| | | KernelVersion | Version number of the operating system | |
| | | SwapSpaceSize | Total swap space size | |
| mw_dba.zen_windowsOS | Retrieves the Windows operating system information | Type | Operating system of the inventoried server | Yes |
| | | OtherTypeDescription | Additional description of the operating system if available | |
| | | Caption | Operating system name | |
| | | CodePage | Language code page of the operating system | |
| | | Version | Version number of the operating system | |
| | | InstallDate | Install date of the operating system | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|----------------------|--|-------------------------|---|------------------------|
| | | VirtualMemorySize | Total number of bytes in the virtual address space of the calling process | |
| | | VisibleMemorySize | Total memory as reported by the operating system | |
| | | ProviderName | Name of the provider | |
| mw_dba.zen_NetWareOS | Retrieves the NetWare operating system information | Type | Operating system of the inventoried server | Yes |
| | | Caption | Operating system name | |
| | | CodePage | Language code page of the operating system | |
| | | Version | Version number of the operating system | |
| | | InstallDate | Install date of the operating system | |
| | | VirtualMemorySize | Total number of bytes in the virtual address space of the calling process | |
| | | VisibleMemorySize | Total memory as reported by the operating system | |
| | | SizeStoredInPagingFiles | The total number of KBytes that can be stored in the OperatingSystem's paging files | |
| | | ProviderName | Name of the provider | |
| | | AccountingVersion | NetWare server specific attributes | |
| | | InternetBridgeSupport | NetWare server specific attributes | |
| | | MaxNumberOfConnections | NetWare server specific attributes | |
| | | MaxNumberOfVolumes | NetWare server specific attributes | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|---------------------|------------------------------------|--------------------------|---|------------------------|
| | | PeakConnectionsUsed | NetWare server specific attributes | |
| | | PrintServerVersion | NetWare server specific attributes | |
| | | QueuingVersion | NetWare server specific attributes | |
| | | RevisionLevel | NetWare server specific attributes | |
| | | SecurityRestrictionLevel | NetWare server specific attributes | |
| | | SFTLevel | NetWare server specific attributes | |
| | | TTSlevel | NetWare server specific attributes | |
| | | VAPVersion | NetWare server specific attributes | |
| | | VirtualConsoleVersion | NetWare server specific attributes | |
| | | InternalNetworkNumber | NetWare server specific attributes | |
| mw_dba.zen_software | Retrieves the software information | Name | Vendor-defined name of the product represented as a vendor trademark or registered trademark. | Yes |
| | | Vendor | Vendor name of the software | |
| | | Version | User-friendly version of a product. For example, the version for Windows 2000 is 2000 or Major.Minor Version of the Product. | |
| | | ProductID | A unique, 16-character identifier for an installed product. This identifier is available from MSI on Windows The format is ABCD-1234-WXYZ-PQRS | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|---------------------|---------------|--------------------|--|------------------------|
| | | InternalVersion | Internal version of a product The format is: <i>major version.minor version.build.sub build number</i> or <i>major version.minor version.build</i> | |
| | | Language | User-friendly name for the language of this copy of the product | |
| | | FriendlyName | Display name of the software | |
| | | Uninstallstring | The command to invoke for uninstalling this product instance. Currently, this is available in Add/Remove Programs (ARP) and MSI on Windows | |
| | | Supportpack | Installed support pack number of the product | |
| | | SoftwareEdition | Product edition defined by the vendor. For example, Professional | |
| | | LastExecutionTime | Date and time stamp when the product was last executed | |
| | | Frequencyofusage | Number of times the product is used | |
| | | Description | Description of the product | |
| | | InstallationSource | Identifies the file system path where the installation files were stored when installing this product instance. Currently, this is available in ARP and MSI on Windows | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|--------------------------|--|-----------------------|--|------------------------|
| | | InformationRepository | Source of scan, which can be, Add/Remove Programs, MSI, Software Dictionary, or PRODUCTS.DAT | |
| | | Category | Product category to which the product belongs For example, Office is a part of the Productivity tools category and Solitaire is a game | |
| | | Helplink | Support web site URL for the product that is available in ARP and MSI | |
| | | PackageGUID | Vendor-defined GUID for a product that is available in MSI | |
| | | Path | Directory path where the product is installed on the inventoried server | |
| mw_dba.zen_softwaregroup | Retrieves the software group information | Name | Vendor-defined name of the software group represented as a vendor trademark or registered trademark | Yes |
| | | Vendor | Vendor name for the software group | |
| | | Version | User-friendly version of a software group | |
| | | ProductID | A unique, 16-character identifier for an installed product. This identifier is available from MSI on Windows The format is ABCD-1234-WXYZ-PQRS. | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|---------------------|---------------|--------------------|--|------------------------|
| | | InternalVersion | Internal version of a product The format is: <i>major version.minor version.build.sub build number</i> or <i>major version.minor version.build</i> | |
| | | Language | User-friendly name for the language of this copy of the product | |
| | | FriendlyName | Display name of the software | |
| | | Uninstallstring | The command to invoke for uninstalling this product instance. Currently, this is available in Add-Remove Programs (ARP) and MSI on Windows | |
| | | Supportpack | Installed support pack number of the product | |
| | | SoftwareEdition | Product edition defined by the vendor. For example, Professional. | |
| | | LastExecutionTime | Date and time stamp when the product was last executed | |
| | | Frequencyofusage | Number of times the product group is used | |
| | | Description | Description of the product group | |
| | | InstallationSource | Identifies the file system path where the installation files were stored when installing this product instance. Currently, this is available in ARP and MSI on Windows | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|--------------------------|---|-----------------------|---|------------------------|
| | | InformationRepository | Source of scan, which can be, Add/Remove Programs, MSI, Software Dictionary, or PRODUCTS.DAT | |
| | | Category | Product category to which the product belongs For example, Office is a part of the Productivity tools category and Solitaire is a game | |
| | | Helplink | Support web site URL for the product that is available in ARP and MSI | |
| | | PackageGUID | Vendor-defined GUID for a product that is available in MSI | |
| | | Path | Directory path where the product is installed on the inventoried server | |
| mw_dba.zen_softwarepatch | Retrieves the software patch information | productid | Software ID of the software patch | No |
| | | PatchName | Vendor-defined name for the patch | |
| mw_dba.zen_antivirus | Retrieves the antivirus product information | Name | Vendor-defined name of the antivirus product represented as a vendor trademark or registered trademark | Yes |
| | | Vendor | Vendor name for the antivirus product | |
| | | Version | User-friendly version of the antivirus product | |
| | | ProductID | A unique, 16-character identifier for an installed antivirus product. This identifier is available from MSI on Windows The format is ABCD-1234-WXYZ-PQRS | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|---------------------|---------------|-------------------|--|------------------------|
| | | InternalVersion | Internal version of the antivirus product The format is: <i>major version.minor version.build.sub build number</i> or <i>major version.minor version.build</i> | |
| | | Language | User-friendly name for the language of this copy of antivirus product | |
| | | FriendlyName | Display name of the antivirus product | |
| | | Uninstallstring | The command to invoke for uninstalling this product instance. Currently, this is available in Add/Remove Programs (ARP) and MSI on Windows | |
| | | Supportpack | Installed support pack number of the antivirus product | |
| | | SoftwareEdition | Antivirus Product edition defined by the vendor | |
| | | LastExecutionTime | Date and time stamp when the antivirus product was last executed | |
| | | Frequencyofusage | Number of times the antivirus product is used | |
| | | Description | Description of the antivirus product | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|---------------------------|---|-----------------------|--|------------------------|
| | | InstallationSource | Identifies the file system path where the installation files were stored when installing this antivirus product instance. Currently, this is available in ARP and MSI on Windows | |
| | | InformationRepository | Source of scan, which can be Add/Remove Programs, MSI, Software Dictionary, or PRODUCTS.DAT | |
| | | DefinitionDate | The date of the virus definition file installed on the computer. Some anti-virus products combine date and version into a single string | |
| | | DefinitionVersion | The vendor-defined version of the virus definition file that has been installed on a computer | |
| | | Category | Product category to which the antivirus product belongs | |
| | | Helplink | Support Web site URL for the antivirus product that is available in ARP and MSI | |
| | | PackageGUID | Vendor-defined GUID for the antivirus product that is available in MSI | |
| | | Path | Directory path where the antivirus product is installed on the inventoried server | |
| mw_dba.zen_dictionaryfile | Retrieves the ZENworks software dictionary file information | fileid | Dictionary File ID | Yes |
| | | directoryid | Directory ID | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|----------------------------|---|----------------------|--|-------------------------------|
| | | FileName | Filename of the dictionary file | |
| | | Directory | Directory name in which the dictionary file is stored | |
| | | FileVersion | Dictionary file version | |
| | | FileSize | Dictionary file size | |
| | | LastModified | Last modified date of the dictionary file | |
| | | InternalName | Internal name | |
| | | ProductVersion | The version of the product represented by this file | |
| | | Company | Vendor name | |
| | | ProductName | The product which this file represents | |
| | | Language | User-friendly name for the language of this copy of the file | |
| | | SoftwareDictionaryID | ID of the file as represented in the General software dictionary | |
| mw_dba.zen_excludedfile | Retrieves the excluded file information | fileid | Excluded file ID | Yes |
| | | directoryid | Directory ID | |
| | | FileName | Filename of the excluded file | |
| | | Directory | Directory name in which the excluded file is stored | |
| | | FileVersion | Excluded file version | |
| | | FileSize | Excluded file size | |
| | | LastModified | Last modified date of the excluded file | |
| | | InternalName | Internal name | |
| | | ProductVersion | The version of the product represented by this file | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|--------------------------|--|-------------------|--|------------------------|
| | | Company | Vendor name | |
| | | ProductName | The product which this file represent | |
| | | Language | User-friendly name for the language of this copy of the file | |
| mw_dba.zen_locktable | Retrieves the lock table information | ComputerName | Computer name | Yes |
| | | LastScanTime | The date when the Scanner was last scanned. Stored as milliseconds time value so it can be read and displayed in any appropriate date format | |
| | | RecentInformation | Latest information | |
| mw_dba.zen_removabledisk | Retrieves the removable disk information | Manufacturer | Vendor name for the removable disk | No |
| | | Description | Description of the removable disk | |
| | | PhysicalCylinders | Total number of cylinders or tracks on the disk | |
| | | PhysicalHeads | Number of heads | |
| | | SectorsPerTrack | Number of sectors per track | |
| | | Capacity | Total size | |
| mw_dba.zen_fixeddisk | Retrieves the fixed disk information | Manufacturer | Vendor name of the fixed disk | No |
| | | Description | Description of the fixed disk | |
| | | PhysicalCylinders | Total number of cylinders or tracks on the disk | |
| | | PhysicalHeads | Number of heads | |
| | | SectorsPerTrack | Number of sectors per track | |
| | | Capacity | Total size | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|-----------------------------------|---|--------------------|---|-------------------------------|
| mw_dba.zen_WindowsLocalFileSystem | Retrieves the Windows local file system information | Name | Windows local file system name | No |
| | | FileSystemSize | Windows local file system size | |
| | | AvailableSpace | Windows local file system space | |
| | | FileSystemType | Windows local file system type | |
| | | Caption | Windows local file system caption | |
| | | DeviceID | Device ID | |
| | | VolumeSerialNumber | Windows local file system volume number | |
| mw_dba.zen_NetWareLocalFileSystem | Retrieves the NetWare local file system information | Name | NetWare local file system name | No |
| | | FileSystemSize | NetWare local file system size | |
| | | AvailableSpace | NetWare local file system available space | |
| | | FileSystemType | NetWare local file system type | |
| | | Caption | NetWare local file system caption | |
| | | DeviceID | Device ID | |
| | | VolumeSerialNumber | NetWare local file volume serial number | |
| mw_dba.zen_LinuxLocalFileSystem | Retrieves the Linux local file system information | Name | Linux local file system name | No |
| | | FileSystemSize | Linux local file system size | |
| | | AvailableSpace | Linux local file system available space | |
| | | FileSystemType | Linux local file system type | |

| Inventory View Name | Functionality | Attributes | Description | Is the view Localized? |
|---------------------|---------------|--------------------|--|------------------------|
| | | Caption | Linux local file system caption | |
| | | DeviceID | Device ID | |
| | | VolumeSerialNumber | Linux local file system volume serial number | |

17.3.2 How to Use the Inventory Views

You can use the Inventory views along with SQL statements, and execute the SQL statements from the Inventory database prompt or in any third-party database front-end application.

Examples:

- ◆ To retrieve all the processor information:


```
select * from mw_dba.zen_processor
```
- ◆ To retrieve specific processor information:


```
select DeviceID, Description, Role, Family,
       OtherFamilyDescription, UpgradeMethod, MaxClockSpeed,
       CurrentClockSpeed from mw_dba.zen_processor
```
- ◆ To retrieve all software information:


```
select * from mw_dba.zen_software
```
- ◆ To retrieve software information along with its suite details:


```
select soft.name, softsuite.name from mw_dba.zen_software soft,
       mw_dba.zen_softwaregroup softsuite where soft.name = 'ZENworks
       Desktop Management Inventory Server' and
       soft.parentinstanceid=softsuite.pinstanceid;
```
- ◆ To retrieve software patch information:


```
select suite.name, patchname from mw_dba.zen_softwaregroup suite,
       mw_dba.zen_softwarepatch patch where
       suite.pinstanceid=patch.pinstanceid
```
- ◆ To retrieve software suite patch information:


```
select soft.name, patchname from mw_dba.zen_software soft,
       mw_dba.zen_softwarepatch patch where
       soft.pinstanceid=patch.pinstanceid;
```
- ◆ To retrieve anti-virus software information:


```
select * from mw_dba.zen_antivirus
```


Monitoring Server Inventory Using Status Logs

18

Novell® ZENworks® 7 Server Inventory with lets you track whether the scan or the roll-up of information is successful by viewing the log files for scan status, roll-up status and Inventory server status.

The scan status of the inventoried server is reported through local log files.

The inventory components report the status of the Inventory server and roll-up of scan information in Novell eDirectory™.

For example, when you view the status logs, you can determine whether the processing of the scan files was successful or if there were any errors while scanning the server or at the time of roll-up.

You can view the following status information:

- ♦ [Section 18.1, “Viewing the Scan Status of an Inventoried Server,” on page 711](#)
- ♦ [Section 18.2, “Viewing the Roll-Up History of the Inventory Server,” on page 712](#)
- ♦ [Section 18.3, “Viewing the Status of Inventory Components on an Inventory Server,” on page 712](#)
- ♦ [Section 18.4, “Viewing the Status of the Last Scan on the Inventoried Server,” on page 713](#)
- ♦ [Section 18.5, “Viewing the Roll-Up Log for the Inventory Servers,” on page 713](#)
- ♦ [Section 18.6, “Exporting the Inventory Status Log Files,” on page 714](#)
- ♦ [Section 18.7, “Overview of Status Logs and Scan Logs,” on page 714](#)
- ♦ [Section 18.8, “Viewing the Status Log in XML Format,” on page 715](#)

18.1 Viewing the Scan Status of an Inventoried Server

The Inventory Agent reports status information and errors in the `invagent.log` file. This log file is stored in the `sys:\etc` directory on NetWare® servers and in the `temp` directory or the `windows\temp` directory on Windows servers.

The native scanner reports status information and errors in the `invnative.log` file. This log file is stored in the `sys:\etc` directory on NetWare servers and in the `temp` directory or the `windows\temp` directory on Windows servers.

The Inventory Policy Enforcer writes the status of the current invocation by the policy engine into the `invagentpolicyenforcer.log` file. You can set the debug flag in the file `invsetup.ini` located in `sys:\system` or `%systemroot%`.

In the `forceDebug=true` mode, the Inventory Agent writes the status of the `.str` file transfer into the `invagentstrtransfer.log` file. This file will be located in the `sys:\system\invscan` directory on NetWare servers and in the `%system_drive%\invscan` directory on Windows servers.

18.2 Viewing the Roll-Up History of the Inventory Server

The Roll-Up Status reports the status of the roll-up information from the Inventory server that initiated the roll-up of information. For example, if your inventory setup consists of a Leaf Server that initiates the roll-up of information to the next-level Root Server, the Roll-Up log displays the roll-up history of the Leaf Server.

The inventory components of the Inventory server (Sender, Receiver, and Storer) write the scan information in the Roll-Up Status. For example, you view the Roll-Up log to determine whether there were any errors during roll-up of inventory information from the Inventory server. This log also displays the most recent roll-up time of the inventory information that was stored in the database on the topmost level server (Root Server). This log displays the history of the ten previous roll-up sessions done from the Inventory server.

Table 18-1 lists the details of the log:

Table 18-1 Details available in the Roll-Up log

| Status Information | Details |
|--------------------|---|
| Roll-Up Start Time | Displays the date and time of the roll-up. |
| Message | Displays the message reported by the inventory component while moving the inventory information across the Inventory servers. |

You can export the file as a .csv or tab-delimited file.

To invoke the Roll-Up Status window:

- 1 In ConsoleOne®, right-click the Inventory Service object (Inventory Service *server_name*), from which the roll-up is done, click *Properties*, click *Status Report* tab, then click *Roll-Up Status*.

18.3 Viewing the Status of Inventory Components on an Inventory Server

The Server Status window reports the status of the Inventory server components on the selected Inventory server. You can view the Inventory server Status log for any Inventory Service object. For example, you can determine whether the Sender sent the files to the Receiver or whether the Storer was able to establish the connection with the database successfully. The Server Status window displays the details of the ten latest status messages logged by the Inventory server components.

If the Inventory server components (Sender, Receiver, Selector, Storer, Scan Collector, Service Manager, or Roll-Up Scheduler) are not up and running on the Inventory server, the status of the Inventory server displays the information.

Table 18-2 lists the details of the log:

Table 18-2 Inventory Details displayed in the Server Status window

| Status Information | Details |
|--------------------|---|
| Time of Log | Displays the date and time when the message was reported by the inventory components. |
| Source | Displays the inventory component that has logged the status message. |
| Message Type | Displays the severity of the message. |
| Message | Displays the message reported by the inventory components. |

You can export the log file as a `.csv` or tab-delimited file.

To view the Server Status window:

- 1 In ConsoleOne, right-click the Inventory Service object (Inventory Service_*server_name*), then click *Properties*, click *Status Report*, then click *Server Status*.

18.4 Viewing the Status of the Last Scan on the Inventoried Server

On NetWare, Windows servers, the `invagent.log` and the `invnative.log` files store the details and last execution status of the Inventory scan.

18.5 Viewing the Roll-Up Log for the Inventory Servers

The Roll-Up log reports the status of the latest roll-up from the Inventory Service objects in the container. For example, you view the Roll-Up log to determine whether the latest roll-up of information from the Roll-Up server for the Inventory Service object was successful. The inventory components (Sender, Receiver, and Storer) write the roll-up information in the Roll-Up log. You can also choose to display error, warning, and informational status messages of the Intermediate servers.

Table 18-3 lists the details of the log:

Table 18-3 Details available in the Roll-Up log

| Status Information | Details |
|------------------------|--|
| Roll-Up Initiated From | Displays the DN of the Intermediate Server that initiated the roll-up. |
| Roll-Up Start Time | Displays the date and time the roll-up of information was initiated. |
| Source | Displays the inventory component that logs the status. |
| Message Type | Displays the severity of the message. |
| Message | Displays the message reported by the inventory components while scanning the inventoried server. |

You can export the log as a `.csv` or tab-delimited file.

To invoke the Roll-Up Log window:

- 1 In ConsoleOne, click the container that contains the Inventory Service object (Inventory Service_server_name), click *Tools*, click *ZENworks Inventory*, then click *Roll-Up Log*.
- 2 Click the severity type of the messages you want to view, then click *OK*.

18.6 Exporting the Inventory Status Log Files

You can store the details of the log files as Comma-Separated-Value reports or as a tab-delimited file.

To save the log as a file:

- 1 In ConsoleOne, open the Status window.
- 2 Click *Export*.
- 3 Select the file type, and specify the filename.
- 4 Click *OK*.

18.7 Overview of Status Logs and Scan Logs

Table 18-4 List of the Inventory status logs and scan logs

| Status/Scan Log | Inventory Components that Log the Status | Details of the Log | How to View the Log File |
|-----------------------------|--|---|--|
| Inventoried Server Scan Log | Scan program, Policy Enforcer | Format module name, time stamp, status code and status message | Available locally on the inventoried server |
| Roll-Up Log | Sender, Receiver, Storer | Roll-up initiated from, roll-up start time, inventory component, message type, status message | Click the container for the Inventory Service object, click <i>Tools</i> , click <i>ZENworks Inventory</i> , then click <i>Roll-Up Log</i> |
| Invagent.log | Scan program, Inventory Agent | Format module name, time stamp, status code and status message | Opens in any text editor |
| Invnativelog | Scan program | Format module name, time stamp, status code and status message | Opens in any text editor |
| Invagentpolicyenforcer.log | Policy Enforcer | Time of log, error type, description, severity and state | Opens in any text editor |

| Status/Scan Log | Inventory Components that Log the Status | Details of the Log | How to View the Log File |
|--|--|--|--|
| Invagentstrtransfer.log (created in the debug mode) | Inventory Agent | Time of log, error type, description, severity and state | Opens in any text editor |
| Status of Inventory components on Server | Sender, Receiver, Scan Collector, Selector, Storer, Service Manager, Roll-Up Scheduler | Time of log, source, message type, message | In ConsoleOne, right-click the Inventory Service object, click <i>Properties</i> , click <i>Status Report</i> , then click <i>Server Status</i> |
| Roll-Up Status | Sender, Receiver, Storer | Roll up start time, message | In ConsoleOne, right-click the Inventory Service object, click <i>Properties</i> , click <i>Status Report</i> , then click <i>Roll-Up Status</i> |

18.8 Viewing the Status Log in XML Format

All inventory components log the status messages in a log file maintained in XML (Extensible Markup Language) format. Unlike the status logs that contain a history of the ten latest status messages, the status XML log stores all status messages.

The log file contains the following information:

- ◆ Inventory module name
- ◆ Date and time of status logging
- ◆ Severity of the message
- ◆ Message text and status message number
- ◆ DN name, if the inventory module is associated with a particular DN object in eDirectory
- ◆ Product-specific details of the module

The format of the log file is as follows:

```
?xml version="1.0" encoding="UTF-8"?>
?xml stylesheet type="text/xsl" href="inventorylog.xsl"?
<message_log>
  <message_entry>
    <module_name>Scanner</module_name>
    <severity>Critical</severity>
    <date_time>8/3/00 12:49 PM</date_time>
    <message_tag>unable to create scan data files
    </message_tag>
    <dn_name>Inv_server</dn_name>
  </message_entry>
</module_name>Storer</module_name>
  <severity>Critical</severity>
```

```
<date_time>8/3/00 12:49 PM</date_time>
<message_tag>unable to update the database</message_tag>
<dn_name>Inv_server</dn_name>
</message_entry>
..
</message_log>
```

A sample style sheet and Document Type Declaration (DTD) file are located in *inventory_installation_directory\inv\server\xmllog* on the Inventory server.

The *inventorylog.xml* log file is located in the *inventory_installation_directory\inv\server\xmllog* directory on NetWare and Windows Inventory servers.

By default, the maximum size of the log file is 100 KB. To modify the maximum size of the log file, edit the *inventorylog.ini* file. On NetWare and Windows Inventory servers, this file is in the *inventory_installation_directory\inv\server\xmllog* directory.

The contents of *inventorylog.ini* are as follows:

```
max_file_size=100 KB
```

Modify the *MAX_FILE_SIZE* parameter, if required.

If the file size exceeds the value specified in the *MAX_FILE_SIZE* parameter, the file is archived as *filename_old.xml*. The latest messages are in the current log file.

To view the log data file, use a third-party XML browser.

Performance Tips

J

This section provides information on the system and database parameters that you need to tune to obtain improved performance for the Server Inventory component of Novell® ZENworks® 7 Server Management. Specific tuning tips are provided for working with Inventory Reports, Database Export, and Query.

In addition to reviewing this information, we recommend that you refer to vendor documentation or other related articles regarding performance tuning and database tuning available on the Internet

This chapter contains the following sections:

- ♦ [Section J.1, “Database Parameter Tuning Tips,” on page 717](#)
- ♦ [Section J.2, “Improving the Throughput of the Inventory Storer,” on page 721](#)
- ♦ [Section J.3, “Performance Tips for the Inventory Server \(Support Pack 1\),” on page 724](#)
- ♦ [Section J.4, “Performance Tips for the Inventory ConsoleOne Utilities,” on page 726](#)
- ♦ [Section J.5, “References,” on page 727](#)

J.1 Database Parameter Tuning Tips

- ♦ [Section J.1.1, “Sybase in the NetWare and Windows Environments,” on page 717](#)
- ♦ [Section J.1.2, “Oracle in the NetWare, Windows, and Linux Environments,” on page 719](#)
- ♦ [Section J.1.3, “Optimizing the Performance of the Oracle Database,” on page 720](#)
- ♦ [Section J.1.4, “MS SQL in the Windows Environment,” on page 721](#)

J.1.1 Sybase in the NetWare and Windows Environments

- ♦ We recommend you to set the database cache size as follows by configuring the `-c` parameter in the Sybase startup:

Table J-1 Recommended total system memory and Sybase cache memory

| Inventoried Servers in the Database (thousands) | Total Memory of the System | Sybase Cache Memory |
|---|----------------------------|---------------------|
| less than 1000 | 384 MB | 128 MB |
| 1 - 5 | 512 MB | 128 MB |
| 5 - 10 | 512 MB - 768 MB | 256 MB |
| 10 - 25 | 768 MB - 1 GB | 256 MB - 400 MB |
| greater than 25 | 1 - 2 GB | 30 - 40% of RAM |

- ♦ If you have more than 5,000 inventoried servers, we recommend that you use multiprocessors for servers hosting the database and span the data files.

- ◆ If you have more than 10,000 inventoried servers, we recommend that you use a dedicated server for the database.
- ◆ Ensure that the drives in which the database files are located have sufficient free disk space for storing the temporary files generated during the operations of Inventory ConsoleOne utilities.
- ◆ If the Storer is taking significant time to store the inventory information in the following scenarios, you can run the Sybindex utility to improve the Storer performance:
 - ◆ Many Inventory agents are simultaneously upgraded to ZENworks 7 and subsequently, all these agents send the full scans for the time to the Inventory server.
 - ◆ The administrator manually triggers full scan from the Inventory Service object resulting in all Inventory agents send the full scan to the Inventory server.
 - ◆ The Inventory database is either re-installed or changed and then the administrator manually triggers full scan from the Inventory Service object resulting in all Inventory agents sending the full scan to the Inventory server.

Before running the Sybindex utility, make sure that the Sybase Inventory database is up and running, and then stop the Storer.

If you have ZENworks 7 Server Management installed, do the following to run the Sybindex utility. If you have ZENworks 7 Server Management with Support Pack 1 installed, see [Section J.3, “Performance Tips for the Inventory Server \(Support Pack 1\),” on page 724](#) to run the Sybindex utility.

On a NetWare server: At the server console prompt, enter `sybindex`.

On a Windows server: At the server command prompt, go to `inventory_server_installation_path\zenworks\inv\server\wminv\bin` and enter `sybindex`.

NOTE: If the Sybase Inventory database is either not hosted on the current Inventory server or is running on a port other than 2638, edit the `sybindex.ncf` (on NetWare), or `sybindex.bat` (on Windows) to change the host and port before running `sybindex`.

For more information:

- ◆ [“Changing the Database Cache Size on a NetWare Database Server” on page 718](#)
- ◆ [“Changing the Database Cache Size on a Windows Database Server” on page 719](#)

Changing the Database Cache Size on a NetWare Database Server

- 1 Stop the Inventory service. For more information, see [Section 13.1.4, “Starting and Stopping the Inventory Service,” on page 482](#).
- 2 Close all connections to the Inventory database.
- 3 Quit the Sybase server.
- 4 Open the `mgmt dbs.ncf` file in the `sys:\system` directory.
- 5 Modify the `-c` parameter.
For example, `-c 64M` sets the cache size to 64 MB.
- 6 Save the file.
- 7 On the server console, load the Inventory database. Enter `MGMTDBS`.
- 8 Start the Inventory service. For more information, see [Section 13.1.4, “Starting and Stopping the Inventory Service,” on page 482](#).

Changing the Database Cache Size on a Windows Database Server

- 1 Stop the Inventory service. For more information, see [Section 13.1.4, “Starting and Stopping the Inventory Service,” on page 482.](#)
- 2 Stop the Sybase service.
On Windows 2000/2003, in the Control Panel, double-click *Administrative Tools*, double-click *Services*, select *Novell Database - Sybase*, then click *Stop*.
- 3 On the database server, run the `ntdbconfig.exe` file from the `inventory_database_installation_path\zenworks\database\dbengine` directory.
`Ntdbconfig.exe` is a ZENworks database configuration utility for the ZENworks database using Sybase on Windows servers. This utility enables you to reconfigure the Sybase service. For the list of parameters recommended by Sybase, see [“Understanding the Sybase Database Startup Parameters” on page 497.](#)
- 4 Modify the `-c` parameter.
- 5 Click OK.
- 6 Restart the Sybase service.
On Windows 2000/2003, in the Control Panel, double-click *Administrative Tools*, double-click *Services*, select *Novell Database - Sybase*, then click *Start*.
- 7 Start the Inventory service. For more information, see [Section 13.1.4, “Starting and Stopping the Inventory Service,” on page 482.](#)

J.1.2 Oracle in the NetWare, Windows, and Linux Environments

- ◆ Use the following memory recommendations:

Table J-2 Recommended total system memory and Oracle SGA memory

| Inventoried Servers in the Database (thousands) | Total Memory of the System | Oracle SGA Memory |
|---|----------------------------|-------------------------|
| less than 1 | 512 MB | 128 MB |
| 1 - 5 | 768 MB | 256 MB |
| 5 - 10 | 1 GB | 400 MB |
| 10 - 100 | 1 GB - 2 GB | 40% of the total memory |

- ◆ Stop unnecessary services and applications running on the server to enable a background service such as Oracle server to run.
- ◆ Ensure that the drives in which the database files are located have sufficient free disk space for storing the temporary files generated during the operations of Inventory ConsoleOne utilities.
- ◆ We recommend that you use a dedicated server to host the Oracle database.
- ◆ Span the data files across multiple physical disks if you have more than 10,000 inventoried servers.

- ◆ Set the virtual memory value between 2 - 4 times the RAM.
- ◆ We recommend that you use multiprocessors for servers hosting the database.
- ◆ Refer to the Oracle performance tuning documentation and other general recommendations that are listed in the [Section J.5, “References,” on page 727](#) section.
 - ◆ Reduce the priority of the foreground application.
 - ◆ Reduce the file cache value and maximize data for network applications.
- ◆ Modify the `init.ora` file for specific organizational requirements.

For example, to obtain about 260 MB of Oracle SGA with `db_block_size=4096`, modify the `init.ora` file with the following values:

```
db_block_buffers = 50000
shared_pool_size = 32768000
sort_area_size = 10000000
```

- ◆ Invoke and append lines to the `_start.sql` file. The `_start.sql` file is invoked by the `mgmtdbo.ncf` or the `mgmtdbo.bat` file when you start the Inventory database instance. Append the following lines to the existing `_start.sql` file:


```
connect mw_dba;alter table cim.t$product cache;
```
- ◆ If you run the Inventory database on Oracle9i, you can set `db_cache_size` instead of `db_block_buffers * db_block_size`.
- ◆ If the Storer is taking significant time to store the inventory information during the full scan, run the following script to create an additional index on `InstalledFile` table to improve the storing time:


```
create index i$installedfile$compid on
mw_dba.installedfile(computerid) tablespace cim5 pctfree 0;
```
- ◆ Refer to the Oracle Administration guide or Performance guide for more information.

J.1.3 Optimizing the Performance of the Oracle Database

If you have an Inventory database on Oracle, you can improve the performance of the database when you generate the inventory reports or query the database.

You use the database buffer cache to store the most recently used data blocks. The database cache is determined as `db_block_buffers * db_block_size`. These parameters are specified in the `zenworks\database\init.ora` file on the database server.

`DB_BLOCK_BUFFERS` specifies the number of database buffers. `DB_BLOCK_SIZE` specifies the size of each database buffer in bytes.

The size of each buffer in the buffer cache is equal to the size of the data block.

If there is additional memory, you configure the database cache size by increasing the `DB_BLOCK_BUFFERS` parameter in the `init.ora` file. If you run Inventory database on Oracle9i, you can set `db_cache_size` instead of `db_block_buffers * db_block_size`

For more information for Performing tips, see [Section J.1, “Database Parameter Tuning Tips,” on page 717](#).

J.1.4 MS SQL in the Windows Environment

- ◆ We recommend that you use a dedicated server for MS SQL.
- ◆ On the MS SQL server, ensure that the tempdb system database is located on the drive having sufficient disk space.
- ◆ Boost the MS SQL server priority.
- ◆ Enable optimization for background services.
- ◆ Use the configuration in the following table:

Table J-3 Recommended total system memory, processor speed and MS SQL cache memory

| Inventoried Servers in the Database (thousands) | Total Memory of the System | MS SQL Cache Memory | Processor Speed |
|---|----------------------------|---------------------|-----------------------|
| less than 10 | 512 MB | 256 MB | Pentium III: 450 M Hz |
| 10 - 20 | 512 MB - 1 GB | 256 MB - 384 MB | Pentium 4: 1.8 G Hz |
| 20 - 50 | 1 GB - 1.5 GB | 512 MB - 768 MB | Pentium 4: 1.8 G Hz |

- ◆ Span the data files across the multiple physical disks if you have more than 5,000 inventoried servers.
- ◆ We recommend that you use multiprocessors for servers hosting the database.
- ◆ For additional tips on MS SQL, refer to the [MS SQL Server documentation \(http://www.sql-server-performance.com/default.asp\)](http://www.sql-server-performance.com/default.asp).

J.2 Improving the Throughput of the Inventory Storer

You can now improve the throughput of the Inventory Storer by deploying multiple Root Servers to directly store the inventory data to the Oracle 9.2.0.6 Inventory database.

The following sections provide more information:

- ◆ [Section J.2.1, “Factors to be Considered Before Deployment,” on page 721](#)
- ◆ [Section J.2.2, “Procedure to Improve the Throughput,” on page 722](#)
- ◆ [Section J.2.3, “Recommendations for Administering the ZENworks Inventory Server,” on page 724](#)
- ◆ [Section J.2.4, “Recommendations for Administering the Inventory Database,” on page 724](#)

J.2.1 Factors to be Considered Before Deployment

- ◆ **Network Topology:** The Root Servers and the Database server must be located in the same LAN.
- ◆ **Frequency of Scans Received by the Inventory Server:** Large number of scans to be processed within a short duration. For example, scanning 25000 servers every day.

- ◆ **Scan Type:** An initial FULL scan storage would take more time compared to subsequent DELTA scan times.
- ◆ **Total number of Root Servers:** If the size of the scan files is smaller, you can achieve a better throughput by deploying a maximum of 6 to 8 Root Servers. But if you deploy more than 8 servers, the throughput might degrade. All servers that you plan to deploy must receive approximately equivalent number of scans.

J.2.2 Procedure to Improve the Throughput

- 1 Stop the Inventory service and the Inventory database.
- 2 Configure a minimum of two Root Servers but a maximum of eight Root Servers to store the inventory data to an Oracle 9.2.0.6 Inventory database.
- 3 Ensure that the Database server has the following requirements:
 - ◆ Three physical disks
 - ◆ Each disk has a drive with at least 30 GB free disk space
 - ◆ Two Pentium IV processors with 2.4 GHz and 2 GB RAM

For example, on Windows assume that the C drive is on disk1, the E drive on disk2, and the F drive on disk3. And the F drive contains the database files.

- 4 Create the following directory structure for database files on all the three drives:

`drive_name\zenworks\inventory\oracle\database`

For example:

```
c:\zenworks\inventory\oracle\database\  
f:\zenworks\inventory\oracle\database\  
e:\zenworks\inventory\oracle\database\  

```

Let's assume that all the inventory database files are present in

`f:\zenworks\inventory\oracle\database.`

- 5 Move the following database files from `f:\zenworks\inventory\oracle\database` as explained below:
 - ◆ Move `log1.ora`, `cim8.ora`, `cim81.ora`, `cim82.ora` and `index1.ora` to `c:\zenworks\inventory\oracle\database.`
 - ◆ Move the following files to `e:\zenworks\inventory\oracle\database:`

```
rbs1.ora  
tmp1.ora  
cim1.ora  
cim2.ora  
cim21.ora  
cim3.ora  
cim4.ora  
cim5.ora  
cim51.ora  
cim7.ora  
cim71.ora  
cim72.ora
```

INDEX2.ORA

- 6 Edit the `f:\zenworks\inventory\oracle\database\init.ora` file to set values for the following parameters as mentioned:

```
db_cache_size=700000000 or above
shared_pool_size = 300000000 or above
pga_aggregate_target=300000000 or above
sort_area_size=10000000 or above
log_buffer = 1024000 or above
compatible=8.1.6.0.0 or above
open_cursors=2048
session_cached_cursors=2048
processes=200
```

- 7 Extract the platform-specific `atlasperf_alterctrl.sql` from `ZENworks_installation_directory\zenworks\inv\server\wminv\properties\sql.zip`.

If Oracle is running on Windows, extract `atlasperf_alterctrl.sql` from the `oracle\winntspecific` directory within `sql.zip`.

If Oracle is running on Unix, extract `atlasperf_alterctrl.sql` from the `oracle\unixspecific` directory within `sql.zip`.

- 8 Modify the file paths in `atlasperf_alterctrl.sql`, if required, and execute `atlasperf_alterctrl.sql` at the SQLPLUS prompt.

- 9 Start the Inventory database.

- 10 Extract the `\oracle\common\atlasperf_alterfreelist.sql` file from `ZENworks_installation_directory\zenworks\inv\server\wminv\properties\sql.zip`, and execute `atlasperf_alterfreelist.sql` at the SQLPLUS prompt.

- 11 Open the Oracle Enterprise Manager console, and ensure that all the indices and primary key constraints of the following tables are set to Degree of the Parallel option - Default, NOLOGGING and Free Lists is 10:

```
zenworks.t$installedproduct
cim.t$product
mw_dba.installedsoftwarepatch
mw_dba.patch
mw_dba."file"
mw_dba.installedfile
```

- 12 On all Inventory servers, edit `ZENworks_installation_directory\zenworks\inv\server\wminv\properties\storerdebug.properties` to set the value of the following parameters as mentioned:

```
filebatchupdate=true
cursorclosedelay=500
```

- 13 Start the Inventory services on all the Inventory servers that are connected to this database.

J.2.3 Recommendations for Administering the ZENworks Inventory Server

- ◆ Avoid or minimize the frequency of importing or removing servers because it would result in many FULL scans.
- ◆ Avoid NDS time out of sync situation because it may trigger FULL scans.
- ◆ Trigger FULL scan on the Inventory service object only if required because it would trigger FULL scans on all inventoried machines connected to the Inventory server.
- ◆ Balance the load of inventory scan and zip files on each server.
- ◆ Stagger the inventory scan and the roll-up schedule.
- ◆ Avoid scheduling too many frequent scans and roll-ups such as daily scans and daily roll-ups.
- ◆ Minimize the scanning of unknown application files and tune the software dictionary. For more information, see [Section 16.3.30, “Base-lining the Software Dictionary Deployment,” on page 629.](#)

J.2.4 Recommendations for Administering the Inventory Database

- ◆ Resize the Oracle SGA parameters appropriately to handle the concurrent updates.
- ◆ Configure appropriate database server hardware requirements such as adding memory, disks.
- ◆ If required, rebuild the indices in the database and scatter them to different tablespaces. Do not have more than one index of the same table on a tablespace.
- ◆ Scatter the data files on multiple physical disks.
- ◆ Apply the standard recommendations as suggested in the Oracle administration or Performance guides.
- ◆ Use a dedicated network between the Inventory server and the Inventory database. For example, 100 MBPS.
- ◆ Add enough rollback segments and properly size them to avoid the ORA-01555 error.
- ◆ If a large number of servers are processed for FULL scan, delete old database and use a new database.

J.3 Performance Tips for the Inventory Server (Support Pack 1)

IMPORTANT: Review this section only if you installed ZENworks 7 Server Management with Support Pack 1

The Server Inventory service might demand high (up to 100%) processor utilization in the following scenarios:

- ◆ Many Inventory agents are simultaneously upgraded to ZENworks 7 and subsequently, all these agents send the full scans for the time to the Inventory server.
- ◆ The administrator manually triggers full scan from the Inventory Service object resulting in all Inventory agents send the full scan to the Inventory server.

- ♦ The Inventory database is either re-installed or changed and then the administrator manually triggers full scan from the Inventory Service object resulting in all Inventory agents sending the full scan to the Inventory server.
- ♦ The Server Inventory process or other applications are running on the ZENworks server.
- ♦ The indexes of the Inventory database might have to be recreated.

If the utilization rate is unacceptable, or if the Inventory Storer takes a considerable amount of time to store the inventory data, perform the following tasks to improve the Inventory server performance:

- 1** Stop the Inventory Service. For more information, see [Section 13.1.4, “Starting and Stopping the Inventory Service,” on page 482](#).
- 2** Close all instances of ConsoleOne that are connected to this database.
- 3** If your Inventory database is running on Sybase, modify the database indexes using the sybindex utility.
 - ♦ To run the sybindex utility on a NetWare server:
 1. On the Inventory server, extract `Inventory_server_installation_path\zenworks\inv\server\wminv\properties\sql.zip` to a temporary directory retaining the directory structure. For example, extract `sql.zip` to `sys:\sql`. The temporary directory contains the Sybase directory.
 2. If the Sybase Inventory database is either not hosted on the Inventory server or is running on a port other than 2638, edit `Inventory_server_installation_path\zenworks\inv\server\wminv\properties\sqlupdater.properties` to specify the host and port on which the Sybase Inventory database is running.
 3. At the Inventory server system console prompt, enter:


```
sybindex -path
complete_path_of_sql.zip_extracted_directory\sybase.
```

 For example, `sybindex -path sys:\sql\sybase` where `sql` is the directory to which `sql.zip` is extracted in Step 1.
 - ♦ To run the sybindex utility on a Windows server:
 1. On the Inventory server, extract `Inventory_server_installation_path\zenworks\inv\server\wminv\properties\sql.zip` to a temporary directory retaining the directory structure. For example, extract `sql.zip` to `c:\sql`. The temporary directory contains the Sybase directory.
 2. If the Sybase Inventory database is either not hosted on the Inventory server or is running on a port other than 2638, edit `Inventory_server_installation_path\zenworks\inv\server\wminv\properties\sqlupdater.properties` to specify the host and port on which the Sybase Inventory database is running.
 3. At the Inventory server command prompt, navigate to `Inventory_server_installation_path\zenworks\inv\server\wminv\bin`, and enter `sybindex -path complete_path_of_sql.zip_extracted_directory\sybase.`
 For example, `sybindex -path c:\sql\sybase` where `sql` is the directory to which `sql.zip` is extracted in Step 1.

NOTE: This execution might take significant amount of time to complete depending on the database size.

- 4 If your Inventory database is running MSSQL database, execute the following scripts available in the MSSQL directory of `Inventory_server_installation_path\zenworks\inv\server\wminv\properties\sql.zip` with appropriate user logins as explained below from the MS SQL Query Analyzer:
 - ♦ Log in as CIM and execute `mssql_perf_cim.sql`.
 - ♦ Log in as ZENworks and execute `mssql_perf_zenworks`.
 - ♦ Log in as ManageWise and execute `mssql_perf_managewise`.
 - ♦ Log in as MW_DBA and execute `mssql_perf_mw_dba`.Ignore any warnings related to DROP statements during the script execution.
- 5 Restart the Inventory Service. For more information, see [Section 13.1.4, “Starting and Stopping the Inventory Service,” on page 482](#).
- 6 Restart the ConsoleOne.
- 7 To improve the throughput of the Storer, you can tune the parameters of the service.
 - 7a Stop the Inventory service. For more information, see [Section 13.1.4, “Starting and Stopping the Inventory Service,” on page 482](#).
 - 7b Run a text editor and navigate to the `Inventory_server_installation_path\zenworks\inv\server\wminv\properties` directory.
 - 7c Open the `storerdebug.properties` file, make the following value change to the uncommented parameter:

```
filebatchupdate=true
```
 - 7d Save the changes and close the file.
 - 7e Restart the Inventory service. For more information, see [Section 13.1.4, “Starting and Stopping the Inventory Service,” on page 482](#).

J.4 Performance Tips for the Inventory ConsoleOne Utilities

This section discusses the performance tips for the Inventory ConsoleOne utilities:

- ♦ [Section J.4.1, “Inventory Reports Performance Tips,” on page 727](#)
- ♦ [Section J.4.2, “Inventory Data Export Performance Tips,” on page 727](#)
- ♦ [Section J.4.3, “Inventory Query Performance Tips,” on page 727](#)

NOTE: If the Inventory ConsoleOne utilities are retrieving data from a large database, you must stop the Storer service before running the utilities. This improves the performance of the ConsoleOne utilities.

If your database is huge, ensure that the machine running the Inventory ConsoleOne utilities has at least 5 GB free hard disk space.

J.4.1 Inventory Reports Performance Tips

If you have more than 1000 inventoried servers in your database, listing all of the subreports consumes time. We recommend that you specify the list of subreports. By doing so, the general performance of the reports is improved.

J.4.2 Inventory Data Export Performance Tips

- ◆ To maximize the performance of Inventory Data Export, you need to enable the filter condition in DBExport. Based on the query you specify, DBExport exports only selected software.
- ◆ During export, deselect the attributes that you do not want to use. To do this, use the DBExport and the Required Attributes Only option.
- ◆ Perform the software export separately. This greatly improves the performance of the Non-Software Export function.

J.4.3 Inventory Query Performance Tips

- ◆ Specify queries using the AND condition in multiple groups to increase performance.
- ◆ Split a complex query with several logical operators into multiple groups separated by a logical operator.
- ◆ If you want to use a complex query, increase the database cache size. For more information on tuning databases, see [Section J.1, “Database Parameter Tuning Tips,” on page 717](#).
- ◆ Save fast, narrowed-down queries for future use.
- ◆ Do not invoke the Inventory Query by connecting to a database over a slow link.
- ◆ If a complex query takes more than 10 minutes to execute over a fast link, you probably do not have any inventoried servers that match the query you specified. The following message is displayed:

```
No Computer system matched the query
```

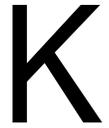
Close the Result window, narrow your input query and retry. Repeat the process of narrowing your query until you locate your inventoried servers.
- ◆ For optimal performance, we recommend that you do not use more than four groups and not more than three logical operators separating the four groups in your query.
- ◆ If you know the exact logical string, avoid using the MATCHES operator. The MATCHES operator searches the database for a result based on the pattern you specify. This results in performance degradation.
- ◆ If you want to check for a particular inventory component not stored in the Inventory database, use the (ISNULL) operator instead of a query with a regular attribute.

J.5 References

For additional information on performance tuning tips, refer to the following documentation for specific components:

- ◆ [MS SQL performance information \(http://www.sql-server-performance.com\)](http://www.sql-server-performance.com)
- ◆ Oracle9i Database and Performance guide and reference
- ◆ Oracle9i Database Administrator's guide

Hardware Information Collected by the Inventory Scanners



This section provides information on the following topics:

- [Section K.1, “Hardware Information Collected on NetWare Inventoried Servers,” on page 729](#)
- [Section K.2, “Hardware Information Collected on Windows Inventoried Servers,” on page 734](#)

K.1 Hardware Information Collected on NetWare Inventoried Servers

Table K-1 Hardware information collected on the NetWare inventoried servers

| Scan Data | SNMP Details | SMBIOS Details |
|-----------------------------|-----------------------|------------------------------|
| System.Type | SNMP v2.0 RFC1213.MIB | Not applicable |
| System.MachineName | SNMP v2.0 RFC1213.MIB | Not applicable |
| System.AssetId | Not applicable | SMBIOS v2.3 Type 3 structure |
| System.Model | Not applicable | SMBIOS v2.3 Type 1 structure |
| System.ModelNumber | Not applicable | SMBIOS v2.3 Type 3 structure |
| System.SystemIdentifier | Not applicable | Not applicable |
| System.ManagementTechnology | Not applicable | Not applicable |
| System.DNSName | Not applicable | Not applicable |
| System.TreeName | Not applicable | Not applicable |
| NetworkAdapter.MACAddress | SNMP v2.0 RFC1213.MIB | Not applicable |
| IP.Address | SNMP v2.0 RFC1213.MIB | Not applicable |
| IP.Subnet (Subnet Mask) | SNMP v2.0 RFC1213.MIB | Not applicable |
| NetworkAdapter.MACAddress | Not applicable | Not applicable |
| IPX.Address | SNMP v2.0 IPX.MIB | Not applicable |
| NetworkAdapter.MACAddress | SNMP v2.0 IPX.MIB | Not applicable |
| DNS.HostName | Not applicable | Not applicable |
| NetworkAdapter.Speed | SNMP v2.0 RFC1213.MIB | Not applicable |
| NetworkAdapter.Name | SNMP v2.0 RFC1213.MIB | Not applicable |
| NetworkAdapter.PermAddress | Not applicable | Not applicable |

| Scan Data | SNMP Details | SMBIOS Details |
|-------------------------------------|-----------------------|------------------------------|
| NetworkAdapter.AdapterType | SNMP v2.0 RFC1213.MIB | Not applicable |
| NetworkAdapter.ProviderName | SNMP v2.0 RFC1213.MIB | Not applicable |
| NetworkAdapter.DriverDescription | SNMP v2.0 RFC1514.MIB | Not applicable |
| NetworkAdapter.DriverName | SNMP v2.0 RFC1514.MIB | Not applicable |
| NetworkAdapter.DriverVersion | SNMP v2.0 RFC1514.MIB | Not applicable |
| Zenworks_ZENNetworkAdapter---offset | SNMP v2.0 RFC1514.MIB | Not applicable |
| Processor.stepping | Not applicable | Not applicable |
| Processor.DeviceID | Not applicable | SMBIOS v2.3 Type 4 structure |
| Processor.Family | Not applicable | SMBIOS v2.3 Type 4 structure |
| Processor.OtherFamily | Not applicable | SMBIOS v2.3 Type 4 structure |
| Processor.MaxClockSpeed | Not applicable | SMBIOS v2.3 Type 4 structure |
| Processor.CurrentClockSpeed | Not applicable | SMBIOS v2.3 Type 4 structure |
| Processor.Role | Not applicable | SMBIOS v2.3 Type 4 structure |
| Processor.UpgradeMethod | Not applicable | SMBIOS v2.3 Type 4 structure |
| Processor.Description | Not applicable | SMBIOS v2.3 Type 4 structure |
| Processor.Name | Not applicable | SMBIOS v2.3 Type 4 structure |
| BIOS.Manufacturer | Not applicable | SMBIOS v2.3 Type 0 structure |
| BIOS.BIOSDate | Not applicable | SMBIOS v2.3 Type 0 structure |
| BIOS.BIOSIDBytes | Not applicable | Not applicable |
| BIOS.Caption | Not applicable | Not applicable |
| BIOS.SerialNumber | Not applicable | Not applicable |
| BIOS.Version | Not applicable | SMBIOS v2.3 Type 0 structure |
| BIOS.PrimaryBIOS | Not applicable | Not applicable |
| BIOS.Size | Not applicable | Not applicable |
| Bus.Type | SNMP v2.0 RFC1514.MIB | Not applicable |
| Bus.Name | Not applicable | Not applicable |

| Scan Data | SNMP Details | SMBIOS Details |
|-------------------------------|-----------------------|-------------------------------|
| Bus.Description | SNMP v2.0 RFC1514.MIB | Not applicable |
| Bus.Version | Not applicable | Not applicable |
| Monitor.NumberOfColorPlanes | Not applicable | Not applicable |
| Monitor.HorizontalResolution | Not applicable | Not applicable |
| Monitor.VerticalResolution | Not applicable | Not applicable |
| Monitor.DisplayType | Not applicable | Not applicable |
| Monitor.MemoryType | Not applicable | Not applicable |
| Monitor.MaxMemorySupported | Not applicable | Not applicable |
| Monitor.Bitsperpixel | Not applicable | Not applicable |
| Monitor.ControllerDescription | Not applicable | SMBIOS v2.3 Type 10 structure |
| Monitor.MaxRefreshrate | Not applicable | Not applicable |
| Monitor.MinRefreshrate | Not applicable | Not applicable |
| Monitor.DACType | Not applicable | Not applicable |
| Monitor.ChipSet | Not applicable | Not applicable |
| Monitor.ProviderName | Not applicable | Not applicable |
| Monitor.VideoBIOSManufacturer | Not applicable | Not applicable |
| Monitor.VideoBIOSVersion | Not applicable | Not applicable |
| Monitor.VideoBIOSReleaseDate | Not applicable | Not applicable |
| Monitor.VideoBIOS.IsShadowed | Not applicable | Not applicable |
| ParallelPort.Name | Not applicable | SMBIOS v2.3 Type 8 structure |
| ParallelPort.DMASupport | Not applicable | Not applicable |
| ParallelPort.Address | Not applicable | Not applicable |
| ParallelPort.IRQ | Not applicable | Not applicable |
| SerialPort.Name | Not applicable | Not applicable |
| SerialPort.Address | Not applicable | SMBIOS v2.3 Type 8 structure |
| SerialPort.IRQ | Not applicable | Not applicable |
| CDROMDrive.DeviceID(*) | Not applicable | Not applicable |
| CDROMDrive.Manufacture | Not applicable | Not applicable |
| CDROMDrive.Description | SNMP v2.0 RFC1514.MIB | Not applicable |
| CDROMDrive.Caption | SNMP v2.0 RFC1514.MIB | Not applicable |
| HardDrive.Media Type | SNMP v2.0 RFC1514.MIB | Not applicable |

| Scan Data | SNMP Details | SMBIOS Details |
|----------------------------------|-----------------------|-------------------------------|
| HardDrive.Vendor | Not applicable | Not applicable |
| HardDisk.Description | SNMP v2.0 RFC1514.MIB | Not applicable |
| HardDisk.Cylinders | Not applicable | Not applicable |
| HardDisk.Heads | Not applicable | Not applicable |
| HardDisk.Sectors | Not applicable | Not applicable |
| HardDisk.Capacity | SNMP v2.0 RFC1514.MIB | Not applicable |
| FileSystem.Name | Not applicable | Not applicable |
| InventoryScanner.Version | Not applicable | Not applicable |
| InventoryScanner.LastScanDate | Not applicable | Not applicable |
| InventoryScanner.InventoryServer | Not applicable | Not applicable |
| InventoryScanner.ScanMode | Not applicable | Not applicable |
| SoundCard.Description | Not applicable | SMBIOS v2.3 Type 10 structure |
| SoundCard.Name | Not applicable | Not applicable |
| SoundCard.Manufacturer | Not applicable | Not applicable |
| Cache.Level | Not applicable | Not applicable |
| Cache.WritePolicy | Not applicable | Not applicable |
| Cache.ErrorCorrection | Not applicable | SMBIOS v2.3 Type 7 structure |
| Cache.Type | Not applicable | SMBIOS v2.3 Type 7 structure |
| Cache.LineSize | Not applicable | Not applicable |
| Cache.ReplacementPolicy | Not applicable | Not applicable |
| Cache.ReadPolicy | Not applicable | Not applicable |
| Cache.Associativity | Not applicable | SMBIOS v2.3 Type 7 structure |
| Cache.Speed | Not applicable | SMBIOS v2.3 Type 7 structure |
| Cache.Size | Not applicable | Not applicable |
| UCS.DNName | Not applicable | Not applicable |
| UCS.PrimaryOwnerContact | Not applicable | Not applicable |
| UCS.PrimaryOwnerName | Not applicable | Not applicable |
| Slot.Description | Not applicable | SMBIOS v2.3 Type 9 structure |

| Scan Data | SNMP Details | SMBIOS Details |
|---|----------------|------------------------------|
| Slot.MaxDataWidth | Not applicable | SMBIOS v2.3 Type 9 structure |
| Slot.ThermalRating | Not applicable | Not applicable |
| LogicalDrive.Name | Not applicable | Not applicable |
| LogicalDrive.DeviceID | Not applicable | Not applicable |
| LogicalDrive.VolumeSerialNumber | Not applicable | Not applicable |
| FileSystem.Name | Not applicable | Not applicable |
| FileSystem.Type | Not applicable | Not applicable |
| FileSystem.TotalSize | Not applicable | Not applicable |
| FileSystem.FreeSpace | Not applicable | Not applicable |
| FileSystem.DeviceID | Not applicable | Not applicable |
| Operating System.OSType | Not applicable | Not applicable |
| OperatingSystem.Version | Not applicable | Not applicable |
| OperatingSystem.Codepage | Not applicable | Not applicable |
| OperatingSystem.InstallDate | Not applicable | Not applicable |
| OperatingSystem.SizeStoredInPagingFiles | Not applicable | Not applicable |
| OperatingSystem.Caption | Not applicable | Not applicable |
| OperatingSystem.TotalVisibleMemorySize | Not applicable | Not applicable |
| OperatingSystem.Role | Not applicable | Not applicable |
| NetWareOperatingSystem.AccountingVersion | Not applicable | Not applicable |
| NetWareOperatingSystem.InternetBridgeSupport | Not applicable | Not applicable |
| NetWareOperatingSystem.MaxNumberOfConnections | Not applicable | Not applicable |
| NetWareOperatingSystem.PeakConnectionsUsed | Not applicable | Not applicable |
| NetWareOperatingSystem.PrintServerVersion | Not applicable | Not applicable |
| NetWareOperatingSystem.QueueingVersion | Not applicable | Not applicable |
| NetWareOperatingSystem.RevisionLevel | Not applicable | Not applicable |
| NetWareOperatingSystem.SecurityRevisionLevel | Not applicable | Not applicable |
| NetWareOperatingSystem.SFTLevel | Not applicable | Not applicable |
| NetWareOperatingSystem.TTSLevel | Not applicable | Not applicable |
| NetWareOperatingSystem.VAPVersion | Not applicable | Not applicable |

| Scan Data | SNMP Details | SMBIOS Details |
|--|----------------|----------------|
| NetWareOperatingSystem.VirtualConsoleVersion | Not applicable | Not applicable |
| NetWareOperatingSystem.InternalNetworkNumber | Not applicable | Not applicable |

K.2 Hardware Information Collected on Windows Inventoried Servers

Table K-2 Hardware information collected on the Windows inventoried servers

| Scan Data | DMI Class and Attribute | WMI Class and Attribute |
|-------------------------------|-----------------------------|---|
| System.Manufacturer | DMTF Component 1 | Win32_ComputerSystemProduct.Vendor |
| System.MachineName | Not applicable | Win32_ComputerSystem.Caption |
| System.AssetTag | DMTF System Enclosure 001.2 | Not applicable |
| System.Model | DMTF Component 2 | Win32_ComputerSystemProduct.Name |
| System.ModelNumber | Not applicable | Not applicable |
| System.SystemIdentifier(GUID) | Not applicable | Not applicable |
| System.SerialNumber | DMTF Component 3 | Win32_ComputerSystemProduct.IndentifyingNumber |
| System.Tag | Not applicable | Not applicable |
| System.ManagementTechnology | Not applicable | Not applicable |
| eDirectory.DNName | Not applicable | Not applicable |
| eDirectory.TreeName | Not applicable | Not applicable |
| NetworkAdapter.MACAddress | Not applicable | Win32_NetworkAdapterConfiguration.MACAddress (Only on Windows 2000; get it through association with Win32_NetworkAdapterSetting) |
| IP.Address | Not applicable | Win32_NetworkAdapterConfiguration.IPAddress (Only on Windows 2000; get it through association with Win32_NetworkAdapterSetting) |

| Scan Data | DMI Class and Attribute | WMI Class and Attribute |
|------------------------------|--|--|
| IP.Subnet (Subnet Mask) | Not applicable | Win32_NetworkAdapterConfiguration.IPSubnet (Only on Windows 2000; get it through association with Win32_NetworkAdapterSetting) |
| NetworkAdapter.MACAddress | Not applicable | Win32_NetworkAdapterConfiguration.MACAddress (Only on Windows 2000; get it through association with Win32_NetworkAdapterSetting) |
| IPX.Address | Not applicable | Win32_NetworkAdapterConfiguration.IPXAddress (Only on Windows 2000; get it through association with Win32_NetworkAdapterSetting) |
| NetworkAdapter.MACAddress | Not applicable | Win32_NetworkAdapterConfiguration.MACAddress (Only on Windows 2000; get it through association with Win32_NetworkAdapterSetting) |
| DNS.HostName | Not applicable | Win32_NetworkAdapterConfiguration.DNSHostName + DNSDomain (Only on Windows 2000; get it through association with Win32_NetworkAdapterSetting) |
| Modem.Description | Not applicable | Win32_POTSModem.Description |
| Modem.Name | Not applicable | Win32_POTSModem.Name |
| Modem.Vendor | Not applicable | Not applicable |
| Modem.DeviceID | Not applicable | Win32_POTSModem.DeviceID |
| NetworkAdapter.DriverVersion | DMTF Network Adapter Driver 001.Driver Software Version | Not applicable |
| Login.CurrentLoggedInUser | Not applicable | Not applicable |
| Login.LastLoggedIn User | Not applicable | Not applicable |
| Login.DomainName | Not applicable | Win32_ComputerSystem.Domain |
| NWClient.Version | Not applicable | Not applicable |
| Processor.stepping | Not applicable | CIM_Processor.Stepping |
| Processor.DeviceID | Not applicable | CIM_Processor.DeviceID |

| Scan Data | DMI Class and Attribute | WMI Class and Attribute |
|-----------------------------|---|--------------------------------------|
| Processor.Family | DMTF Processor 004.3 | CIM_Processor.Family |
| Processor.OtherFamily | Not applicable | CIM_Processor.OtherFamilyDescription |
| Processor.MaxClockSpeed | DMTF Processor 004.5 | CIM_Processor.MaxClockSpeed |
| Processor.CurrentClockSpeed | DMTF Processor 004.6 | CIM_Processor.CurrentClockSpeed |
| Processor.Role | DMTF Processor 004.2 | CIM_Processor.ProcessorType |
| Processor.Upgrade | DMTF Processor 004.7 | CIM_Processor.UpgradeMethod |
| Processor.Description | DMTF Processor 004.4 | CIM_Processor.Description |
| Processor.Name | Enum equivalent of DMTF Processor 004.3 | CIM_Processor.Name |
| BIOS.Manufacturer | DMTF SystemBIOS 001.2 | Win32_BIOS.Manufacturer |
| BIOS.BIOSDate | Not applicable | Win32_BIOS.InstallDate |
| BIOS.BIOSIDBytes | Not applicable | Not applicable |
| BIOS.Copyright | Not applicable | Win32_BIOS.Caption |
| BIOS.SerialNumber | Not applicable | Win32_BIOS.SerialNumber |
| BIOS.BIOSType | DMTF SystemBIOS 001.3 | Win32_BIOS.SMBIOSBIOSVersion |
| BIOS.PrimaryBIOS | DMTF SystemBIOS 001.9 | Win32_BIOS.PrimaryBIOS |
| BIOS.Size | DMTF SystemBIOS 001.4 | Not applicable |
| Bus.Type | Not applicable | Win32_Bus.BusType |
| Bus.Name | Not applicable | Win32_Bus.Name |
| Bus.Description | Not applicable | Win32_Bus.Descriptpion |
| Bus.Version | Not applicable | Not applicable |
| Bus.DeviceID | Not applicable | Win32_Bus.DeviceID |
| IRQ.Number | DMTF IRQ 002.IRQNumber | CIM_IRQ.IRQNumber |
| IRQ.Availability | DMTF IRQ 002.Availablity | CIM_IRQ.Availability |

| Scan Data | DMI Class and Attribute | WMI Class and Attribute |
|--|---|---|
| IRQ.TriggerType | DMTF IRQ 002.TriggerType | CIM_IRQ.TriggerType |
| IRQ.Shareable | DMTF IRQ 002.Shareable | CIM_IRQ.Shareable |
| Keyboard.Layout | DMTF Keyboard 003.Layout | CIM_Keyboard.Layout |
| Keyboard.Subtype | Not applicable | Not applicable |
| Keyboard.Type | DMTF Keyboard 003.Keyboard.Type | CIM_Keyboard.Description |
| Keyboard.Fkeys | Not applicable | CIM_Keyboard.NumberOfFunctionKeys |
| Keyboard.Delay | Not applicable | Not applicable |
| Keyboard.TypeomaticRate | Not applicable | Not applicable |
| VideoAdapter.NumberOfColorPlanes (NEW) | Not applicable | Win32_VideoController.NumberOfColorPlanes |
| VideoAdapter.HorizontalResolution | DMTF Video 004.Current Horizontal Resolution | Win32_VideoController.CurrentHorizontalResolution |
| VideoAdapter.VerticalResolution | DMTF Video 004.Current Vertical Resolution | Win32_VideoController.CurrentVerticalResolution |
| VideoAdapter.DisplayType | DMTF Video 004.Video Type | Win32_VideoController.VideoArchitecture |
| VideoAdapter.MemoryType | DMTF Video 004.Video Memory Type | Win32_VideoController.VideoMemoryType |
| VideoAdapter.MaxMemorySupported | DMTF Video 004.Video RAM Memory Size | Win32_VideoController.AdapterRAM |
| VideoAdapter.Bitsperpixel | DMTF Video 004.Current Number of Bits per Pixel | Win32_VideoController.CurrentBitsPerPixel |
| VideoAdapter.ControllerDescription | DMTF Video 004.Video Controller Description | Win32_VideoController.Description |
| VideoAdapter.MaxRefreshrate | DMTF Video 004.Maximum Refresh Rate | Win32_VideoController.MaxRefreshRate |
| VideoAdapter.MinRefreshrate | DMTF Video 004.Minimum Refresh Rate | Win32_VideoController.MinRefreshRate |
| VideoAdapter.DACType | Not applicable | Win32_VideoController.AdapterDACType |
| VideoAdapter.ChipSet | Not applicable | Win32_VideoController.VideoProcessor |

| Scan Data | DMI Class and Attribute | WMI Class and Attribute |
|---------------------------------|---|---|
| VideoAdapter.ProviderName | Not applicable | Win32_VideoController.VideoAdapterCompatibility |
| VideoBIOS.VideoBIOSManufacturer | DMTF Video BIOS 001.BIOS Manufacturer | CIM_VideoBIOSElement.Manufacturer |
| VideoBIOS.VideoBIOSVersion | DMTF Video BIOS 001.Video.BIOS Version | CIM_VideoBIOSElement.Version |
| VideoBIOS.VideoBIOSReleaseDate | DMTF Video BIOS 001.Video.BIOS Release Date | CIM_VideoBIOSElement.InstallDate |
| VideoBIOS.VideoBIOS.IsShadowed | DMTF Video BIOS 001.Video.Shadowing State | CIM_VideoBIOSElement.IsShadowed |
| ParallelPort.Name | DMTF Parallel Ports 003.Parallel Port Index | CIM_ParallelController.Name |
| ParallelPort.DMASupport | DMTF Parallel Ports 003.DMA Support | CIM_ParallelController.DMASupport |
| ParallelPort.Address | DMTF Parallel Ports 003.Parallel Base I/O Address | Not applicable |
| ParallelPort.IRQ | DMTF Parallel Ports 003.IRQ Used | Not applicable |
| SerialPort.Name | DMTF Serial Ports 004.Serial Port Index | CIM_SerialController.Name |
| SerialPort.Address | DMTF Serial Ports 004.Serial Base I/O Address | Not applicable |
| SerialPort.IRQ | DMTF Serial Ports 004.IRQ Used | Not applicable |
| FloppyDrive.DeviceID | DMTF Logical Drives 001.Logical Drive Name (when DMTF Logical Drives 001.Logical Drive Type=Floppy Drive(7)) | Win32_LogicalDisk.DeviceID (where Win32_LogicalDisk.DriveType = 2 (Removable Disk) and Win32_LogicalDisk.MediaType = [1,10]) |
| FloppyDrive.Manufacture | Not applicable | Not applicable |

| Scan Data | DMI Class and Attribute | WMI Class and Attribute |
|----------------------------------|---|--|
| FloppyDrive.Description | Hard Code: Floppy Drive (when DMTF Disks 003.Storage Type=Floppy Disk(4)) | Win32_LogicalDisk.Description (where Win32_LogicalDisk.DriveType = 2 (Removable Disk) and Win32_LogicalDisk.MediaType = [1,10]) |
| FloppyDrive.MaxNumberOfCylinders | Not applicable | Not applicable |
| FloppyDrive.NumberOfHeads | Not applicable | Not applicable |
| FloppyDrive.SectorsPerTrack | Not applicable | Not applicable |
| FloppyDrive.Size | DMTF Logical Drives 001.Logical Drive Size (when DMTF Logical Drives 001.Logical Drive Type = Floppy Drive(7)) | Win32_LogicalDisk.Size (where Win32_LogicalDisk.DriveType = 2 (Removable Disk) and Win32_LogicalDisk.MediaType = [1,10]) |
| CDROMDrive.DeviceID | DMTF Logical Drives 001.Logical Drive Name (When DMTF Logical Drives 001.Logical Drive Type = 6) | Win32_CDROMDrive.Drive |
| CDROMDrive.Manufacture | Not applicable | Win32_CDROMDrive.Manufacturer |
| CDROMDrive.Description | Not applicable | Win32_CDROMDrive.Description |
| CDROMDrive.Caption | Hard code: CDROM Device (when DMTF Disks 001.Logical Drive Type = 6) | Win32_CDROMDrive.Caption |
| HardDrive.Media Type | DMTF Disks 003.Removable Media | Win32_DiskDrive.MediaType |
| HardDrive.Vendor | Not applicable | Win32_DiskDrive.Manufacturer |
| HardDisk.Description | DMTF Disks 003.Interface Description (when DMTF Disks 003.Storage Type=Hard Disk(3)) | Win32_DiskDrive.Description |

| Scan Data | DMI Class and Attribute | WMI Class and Attribute |
|--------------------------------------|--|---|
| HardDisk.Cylinders | DMTF Disks 003.Number of Physical Cylinders | Win32_DiskDrive.TotalCylinders |
| HardDisk.Heads | DMTF Disks 003.Number of Physical Heads | Win32_DiskDrive.TotalHeads |
| HardDisk.Sectors | DMTF Disks 003.Number of Physical Sectors per Track | Win32_DiskDrive.SectorsPerTrack |
| HardDisk.Capacity | DMTF Disks 003.Total Physical Size | Win32_DiskDrive.Size |
| LogicalDrive.Name | Not applicable | Win32_LogicalDiskDeviceID (when Win32_LogicalDisk.DriveType = 3 (Local Disk)) |
| LogicalDrive.VolumeSerialNumber | Not applicable | Win32_LogicalDisk.VolumeSerialNumber (when Win32_LogicalDisk.DriveType = 3 (Local Disk)) |
| LogicalDrive.Volume (Volume Label) | Not applicable | Win32_LogicalDisk.VolumeName (when Win32_LogicalDisk.DriveType = 3 (Local Disk)) |
| OperatingSystem.OSType | Not applicable | Win32_OperatingSystem.OSType |
| OperatingSystem.Version | Not applicable | Not applicable |
| OperatingSystem.Codepage | Not applicable | Win32_OperatingSystem.CodeSet |
| OperatingSystem.InstallDate | Not applicable | Win32_OperatingSystem.InstallDate |
| OperatingSystem.TotalSwapSpaceSize | DMTF System Memory Settings 001.Total Size of Paging Files | Win32_OperatingSystem.SizeStoredInPagingFiles |
| OperatingSystem.Description | DMTF Operating System 001.Operating System Description | Win32_OperatingSystem.Caption |
| OperatingSystem.OtherTypeDescription | Not applicable | Win32_OperatingSystem.OtherTypeDescription |
| OperatingSystem.VirtualMemorySize | DMTF System Memory Settings 001.Total Virtual Memory | Win32_OperatingSystem.TotalVirtualMemory |
| OperatingSystem.VisibleMemorySize | Not applicable | Win32_OperatingSystem.TotalVisibleMemorySize |
| OperatingSystem.Role | Not applicable | Not applicable |

| Scan Data | DMI Class and Attribute | WMI Class and Attribute |
|--|--|-------------------------------------|
| InventoryScanner.Version | Not applicable | Not applicable |
| InventoryScanner.LastScanDate | Not applicable | Not applicable |
| InventoryScanner.InventoryServer | Not applicable | Not applicable |
| InventoryScanner.ScanMode | Not applicable | Not applicable |
| InventoryScanner.GeneralDictionary Version | Not applicable | Not applicable |
| InventoryScanner.PrivateDictionaryVersion | Not applicable | Not applicable |
| SoundCard.Description | Not applicable | Win32_SoundDevice.Description |
| SoundCard.Name | Not applicable | Win32_SoundDevice.Name |
| SoundCard.Manufacturer | Not applicable | Win32_SoundDevice.Manufacturer |
| Cache.Level | DMTF System Cache 003.System Cache Level | Win32_CacheMemory.Level |
| Cache.WritePolicy | DMTF System Cache 003.System Cache Write Policy | Win32_CacheMemory.WritePolicy |
| Cache.ErrorCorrection | DMTF System Cache 003.System Cache Error Correction | Win32_CacheMemory.ErrorCorrectType |
| Cache.Type | DMTF System Cache 003.System Cache Type | Win32_CacheMemory.CacheType |
| Cache.LineSize | DMTF System Cache 003.Line Size | Win32_CacheMemory.LineSize |
| Cache.ReplacementPolicy | DMTF System Cache 003.Replace ment Policy | Win32_CacheMemory.ReplacementPolicy |
| Cache.ReadPolicy | DMTF System Cache 003.Read Policy | Win32_CacheMemory.ReadPolicy |
| Cache.Associativity | DMTF System Cache 003.Associati vity | Win32_CacheMemory.Associativity |
| Cache.Speed | DMTF System Cache 003.System Cache Speed | Win32_CacheMemory.CacheSpeed |
| Cache.Size | DMTF System Cache 003.System Cache Size | Win32_CacheMemory.MaxCacheSize |

| Scan Data | DMI Class and Attribute | WMI Class and Attribute |
|-------------------------------------|---|--|
| MotherBoard.Version | Not applicable | Win32_BaseBoard.Version |
| MotherBoard.Description | Not applicable | Win32_BaseBoard.Description |
| MotherBoard.Slots | DMTF Motherboard 001.Number of Expansion slots | Not applicable |
| MotherBoard.Manufacture | Not applicable | Win32_BaseBoard.Manufacture |
| Battery.Name | DMTF Portable Battery 002.Portable Battery Device Name | Win32_Battery.Name |
| Battery.Chemistry | DMTF Portable Battery 002.Portable Battery Device Chemistry | Win32_Battery.Chemistry |
| Battery.Capacity | DMTF Portable Battery 002.Portable Battery Design Capacity | Win32_Battery.DesignCapacity |
| Battery.Voltage | DMTF Portable Battery 002.Portable Battery Design Voltage | Win32_Battery.DesignVoltage |
| Battery.Version | DMTF Portable Battery 002.Portable Battery Smart Battery Version | Win32_Battery.SmartBatteryVersion |
| Battery.Manufacturer | DMTF Portable Battery 002.Portable Battery Manufacturer | Win32_PortableBattery.Manufacturer |
| Battery.ManufactureDate | DMTF Portable Battery 002.Portable Battery Manufacturer Date | Win32_Battery.InstallDate |
| Battery.SerialNumber | DMTF Portable Battery 002.Portable Battery Serial Number | Not applicable |
| PowerSupply.InputVoltageDescription | DMTF Power Supply 002.Power Supply Input Voltage Capability Description | CIM_UninterruptiblePowerSupply.Description |

| Scan Data | DMI Class and Attribute | WMI Class and Attribute |
|--------------------------------|---|---|
| PowerSupply.Power | DMTF Power Supply 002.Total Output Power | CIM_UninterruptiblePowerSupply.TotalOutputPower |
| DMA.Number | DMTF DMA 001.DMA Number | CIM_DMA.DMAChannel |
| DMA.Description | DMTF DMA 001.DMA Description | CIM_DMA.Description |
| DMA.Availability | DMTF DMA 001.DMA Channel Availability | CIM_DMA.Availability |
| DMA_BurstMode | DMTF DMA 001.DMA BurstMode | CIM_DMA.BurstMode |
| UCS.DNNName | Not applicable | Not applicable |
| UCS.PrimaryOwnerContact | DMTF General Information 001.3 | CIM_UnitaryComputerSystem.PrimaryOwnerContact |
| UCS.PrimaryOwnerName | DMTF General Information 001.4 | CIM_UnitaryComputerSystem.PrimaryOwnerName |
| PointingDevice.DeviceType | DMTF Pointing Device Pointing Device Type(1) | CIM_PointingDevice.PointingType |
| PointingDevice.Type | DMTF Pointing Device Pointing Device Interface (2) | CIM_PointingDevice.Name |
| PointingDevice.NumberOfButtons | DMTF Pointing Device Pointing Device Buttons (4) | CIM_PointingDevice.NumberOfButtons |
| PointingDevice.DriverName | DMTF Pointing Device Pointing Device Driver Name (6) | Not applicable |
| PointingDevice.DriverVersion | DMTF Pointing Device Pointing Device Driver Version (7) | CIM_PointingDevice.Name |
| PointingDevice.IRQ | DMTF Pointing Device Pointing Device IRQ (3) | Not applicable |
| Slot.Description | DMTF System Slots 003.Description | Win32_SystemSlot. SlotDesignation |
| Slot.MaxDataWidth | DMTF System Slots 003.MaxData Width | Win32_SystemSlot. MaxDataWidth |

| Scan Data | DMI Class and Attribute | WMI Class and Attribute |
|-------------------------------|---|---|
| Slot.ThermalRating | DMTF System Slots 003.Slot Thermal Rating | Win32_SystemSlot. ThermalRating |
| FileSystem.Drive | Not applicable | Win32_LogicalDisk.DeviceID (when Win32_LogicalDisk.DriveType = 3 (Local Disk)) |
| FileSystem.FileSystemSize | Not applicable | Win32_LogicalDisk.Size (when Win32_LogicalDisk.DriveType = 3 (Local Disk)) |
| FileSystem.AvailableSpace | Not applicable | Win32_LogicalDisk.FreeSpace (when Win32_LogicalDisk.DriveType = 3 (Local Disk)) |
| FileSystem.FileSystem | Not applicable | Win32_LogicalDisk.FileSystem (when Win32_LogicalDisk.DriveType = 3 (Local Disk)) |
| Monitor.Device ID | Not applicable | Not applicable |
| Monitor.Description | Not applicable | Not applicable |
| Monitor.Manufacturer Date | Not applicable | Not applicable |
| Monitor.Model ID | Not applicable | Not applicable |
| Monitor.ViewableSize (inches) | Not applicable | Not applicable |
| Monitor.NominalSize (inches) | Not applicable | Not applicable |
| Monitor.Serial Number | Not applicable | Not applicable |
| Monitor.Manufacturer | Not applicable | Not applicable |
| Monitor.Model | Not applicable | Not applicable |
| Chassis.Type (enum) | DMTF Physical Container Global Table 1 | Win32_SystemEnclosure. ChassisTypes |
| Chassis.Manufacturer | DMTF FRU 4 | Win32_SystemEnclosure. Manufacturer |
| Chassis.SerialNumber | DMTF FRU 7 | Win32_SystemEnclosure. SerialNumber |
| Chassis.AssetTag | DMTF Physical Container Global Table 2 | Win32_SystemEnclosure. SMBIOSAssetTag |
| Chassis.Version | Not applicable | Win32_SystemEnclosure. Version |
| Chassis.NumberOfPowerCords | Not applicable | Win32_SystemEnclosure. NumberOfPowerCords |
| Chassis.Tag | Not applicable | Win32_SystemEnclosure. Tag |

NOTE: PCMCIA modems are connected to the computer through the PCMCIA slots on the inventoried servers. The Scanner detects PCMCIA modems that are active on the computer. If you want to know which modem is installed on the computer, use the Windows System Device Manager on the Windows server.

Non-PCMCIA modems are connected to the computer through the external ports. For example, some non-PCMCIA modems are connected through the serial ports. The Scanner detects non-PCMCIA modems that are installed on the computer.

Non-PCMCIA modems might not be active at the time of scanning. Also, these modems might not be connected, although they are configured on the computer. In this case, the Scanner detects the modem and reports the scan information of the modem.

The Inventory scanner reports inventory information for the monitors that are manufactured only after 1997.

ZENworks 7 Server Management Inventory Attributes



Table L-1 lists the Server Inventory attributes that ZENworks 7 Server Management uses.

Each row in the table has:

- ◆ Name of the attribute as displayed in the Inventory Database Export Wizard in ConsoleOne
- ◆ Name of the attribute in the exported .csv file (first row in the .csv file)
- ◆ Inventory database attribute name
- ◆ Type of the attribute in the Inventory database
- ◆ Length of the attribute in the Inventory database
- ◆ Brief description of the attribute

Table L-1 Server Inventory attributes used in ZENworks Server Management

| Export Wizard Attribute Name | Export Attribute Name (Column Heading in the .csv file) | Database Schema Attribute Name | Data Type | Length | Description of the Attribute |
|---------------------------------|---|--|-----------|--------|--|
| General-NDSName-Label | NDSName_LABEL | ManageWise.NDSName.Label | String | 254 | The DN name of the inventoried server registered in eDirectory. |
| SystemInfo.Description | Asset_Description | Zenworks.SystemInfo.Description | String | 254 | Description of the system asset information. |
| SystemInfo.Caption | Asset_Caption | Zenworks.SystemInfo.Caption | String | 64 | Identifying information of the computer. |
| SystemInfo.Tag | Asset_Asset Tag | Zenworks.SystemInfo.Tag | String | 254 | Asset tag number that the ROM-based setup program creates. This is unique to every inventoried server. |
| SystemInfo.ModelNumber | Asset_Model Number | Zenworks.SystemInfo.Model | String | 64 | Model number value for the computer, assigned during manufacture. |
| SystemInfo.SerialNumber | Asset_Serial Number | Zenworks.SystemInfo.SerialNumber | String | 64 | Model serial number value for the computer, assigned during manufacture. |
| SystemInfo.ManagementTechnology | Asset_Management Technology | Zenworks.SystemInfo.ManagementTechnology | Integer | | The management technology available on the computer system. |

| Export Wizard Attribute Name | Export Attribute Name (Column Heading in the .csv file) | Database Schema Attribute Name | Data Type | Length | Description of the Attribute |
|---|---|--|-------------------------------|--------|---|
| CurrentLoginUser.Name | Current Login User.Name | ManageWise."User".Name | String | 254 | User logged in to the Primary eDirectory tree when the inventoried server was scanned. |
| LastLoginUser.Name | Last Login User.Name | ManageWise."User".Name | String | 254 | User most recently logged in to the Primary eDirectory tree through Novell Client when the inventoried server was scanned. |
| Product.Name | Applications_Name | CIM.Product.Name | String | 254 | Name of the software application. |
| Product.Vendor | Applications_Vendor | CIM.Product.Vendor | String | 254 | Name of the software application manufacturer. |
| Product.Version | Applications_Version | CIM.Product.Version | String | 64 | Version of the software application. |
| Product.Location | Applications_Path | CIM.Directory.Location | String | 254 | The product installation path. |
| Product.IndentifyingNumber | Applications_Indentifying Number | CIM.Product.IndentifyingNumber | String | 64 | Microsoft product ID |
| WinOperatingSystem.OSType | Windows_Name | ZENworks.WIN OperatingSystem.OSType | Unsigned Small Integer (enum) | | Operating system name. For example, Windows 2000. See Section M.3, "Enumeration Values for Software-Operating Systems-Windows - Name," on page 768. |
| WinOperatingSystem.Version | Windows_Version | ZENworks.WIN OperatingSystem.Version | String | 254 | Version of the operating system. |
| WinOperatingSystem.Caption | Windows_Caption | ZENworks.WIN OperatingSystem.Caption | String | 64 | Short name of the operating system. For example, Windows 2000. |
| WinOperatingSystem.Role | Windows_Role | ZENworks.WIN OperatingSystem.Role | Integer (enum) | | The role of the computer system. For example, server. |
| WinOperatingSystem.OtherTypeDescription | Windows_Other Description | ZENworks.WIN OperatingSystem.Description | String | 254 | More description about the operating system. |
| WinOperatingSystem.InstallDate | Windows_Install Date | ZENworks.ZEN OperatingSystem.InstallDate | String | 25 | Installation date of the operating system. |
| WinOperatingSystem.CodePage | Windows_Code Page | ZENworks.WIN OperatingSystem.CodePage | String | 254 | Current language code page being used. |

| Export Wizard Attribute Name | Export Attribute Name (Column Heading in the .csv file) | Database Schema Attribute Name | Data Type | Length | Description of the Attribute |
|---|---|---|------------------|--------|--|
| WinOperatingSystem.TotalVisibleMemorySize | Windows_Total Memory (MB) | ZENworks.WIN OperatingSystem.TotalVisibleMemorySize | Integer | | Total memory as reported by the Windows operating system. |
| WinOperatingSystem.TotalVirtualMemorySize | Windows_Total Virtual Memory (MB) | ZENworks.WIN OperatingSystem.TotalVirtualMemorySize | | | Total virtual memory as reported by the Windows operating system. |
| InventoryScanner.Version | Scanner Information_Version | ZENworks.InventoryScanner.Version | String | 64 | Version of the Scanner running on the inventoried server. |
| InventoryScanner.LastScanDate | Scanner Information_Last Scan Date | ZENworks.InventoryScanner.LastScanDate | Unsigned Integer | | The date when the Scanner was last scanned. Stored as milliseconds time value so it can be read and displayed in any appropriate date format. |
| InventoryScanner.InventoryServer | Scanner Information_Inventory Server | ZENworks.InventoryScanner.InventoryServer | String | 254 | Name of the Inventory server to which the scans are sent. It is not the complete DN of the server name. |
| InventoryScanner.ScanMode | Scanner Information_Scan Mode | ZENworks.InventoryScanner.ScanMode | Integer (enum) | | The management technology used by the Scanner, such as WMI or DMI, for scanning the computer system. |
| NetWareClient.Version | Netware Client_Version | ZENworks.NetWareClient.Version | String | 64 | Version of the NetWare client software installed on the inventoried server. |
| NetworkAdapterDriver.Description | Network Adapter Driver_Description | ZENworks.NetworkAdapterDriver.Description | String | 254 | Description of the network adapter driver installed on the inventoried server. For example, IBM 10/100 Ethernet adapter, EN-2420Px Ethernet adapter. |
| NetworkAdapterDriver.Name | Network Adapter Driver_Name | ZENworks.NetworkAdapterDriver.Name | String | 254 | Name of the network adapter driver software installed that corresponds to the adapter. For example, ne2000.sys, pppmac.vxd, and others. |
| NetworkAdapterDriver.Version | Network Adapter Driver_Version | ZENworks.NetworkAdapterDriver.Version | String | 64 | Network adapter driver version. |
| PointingDeviceDeviceDriver.Name | Pointing Device Driver_Name | ZENworks.PointingDeviceDeviceDriver.Name | String | 254 | Name of the mouse driver installed on the inventoried server. |

| Export Wizard Attribute Name | Export Attribute Name (Column Heading in the .csv file) | Database Schema Attribute Name | Data Type | Length | Description of the Attribute |
|-------------------------------------|---|---|------------------------|--------|--|
| PointingDevice.DeviceDriver.Version | Pointing Device Driver_Version | ZENworks.PointingDeviceDeviceDriver.Version | String | 64 | Mouse driver version. |
| PointingDevice.Name | Pointing Device_Name | CIM.PointingDevice.Name | String | 254 | <p>The name of the pointing device, such as Mouse. The string stored in this field will be MOUSE.</p> <p>The CIM.PointingDevice.PointingType field determines the type of the pointing device.</p> <p>The different types of pointing devices are as listed in Section M.7, "Enumeration Values for Hardware-Pointing Device-Name," on page 769.</p> |
| PointingDevice.Numberofbuttons | Pointing Device_Number of Buttons | CIM.PointingDevice.NumberOfButtons | Unsigned Tiny Integer | | The number of buttons used by the pointing device. |
| PointingDevice.IRQNumber | Pointing Device_IRQ Number | CIM.IRQ.IRQNumber | Unsigned Integer | | <p>The IRQ channel on the system to which the Mouse pointing device is attached. This information is stored in an IRQ class and not in the PointingDevice class in the database. For more information on how they are associated, see "Understanding the ZENworks 7 Server Managements Inventory Database Schema" on page 549.</p> |
| PointingDevice.PointingType | Pointing Device_Type | CIM.PointingDevice.PointingType | Integer (enum) | | The pointing device type. |
| ZENKeyboard.Numberoffunction keys | Keyboard_Number of Function Keys | ZENworks.ZENKeyboard.NumberOfFunctionKeys | Unsigned Small Integer | | Number of function keys on keyboard. |
| ZENKeyboard.Layout | Keyboard_Layout | ZENworks.ZENKeyboard.layout | String | 254 | Layout information. For example, US English. |
| ZENKeyboard.SubType | Keyboard_Subtype | ZENworks.ZENKeyboard.SubType | Unsigned Integer | | A number indicating the subtype of the keyboard. |

| Export Wizard Attribute Name | Export Attribute Name (Column Heading in the .csv file) | Database Schema Attribute Name | Data Type | Length | Description of the Attribute |
|--|---|--|-----------------------------------|--------|--|
| ZENKeyboard.Delay | Keyboard_Delay (mSecs) | ZENworks.ZENKeyboard.Delay | Unsigned Integer | | Delay before the repeat of a key. |
| ZENKeyboard.TypeAutomaticRate | Keyboard_TypeAutomatic Rate (mSecs) | ZENworks.ZENKeyboard.TypeAutomatic Rate | Unsigned Integer | | Rate of processing the keys. |
| ZENKeyboard.Description | Keyboard_Description | ZENworks.ZENKeyboard.Description | String | 254 | Keyboard description indicating the type of keyboard. For example, IBM enhanced (101/102 key) keyboard. |
| VideoBIOSElement.Manufacturer | Display_Driver_Manufacturer | CIM.VideoBIOSElement.Manufacturer | String | 254 | Manufacturer of the video BIOS driver installed on the system. |
| VideoBIOSElement.Version | Display_Driver_Version | CIM.VideoBIOSElement.Version | String | 254 | Version of the Video BIOS driver. |
| VideoBIOSElement.Install Date | Display_Driver_Install Date | CIM.VideoBIOSElement.InstallDate | String | 25 | Video BIOS release date. |
| VideoBIOSElement.IsShadowed | Display_Driver_IsShadowed | CIM.VideoBIOSElement.IsShadowed | BIT (Used for Boolean conditions) | | A Boolean condition indicating if the video BIOS supports shadow memory. 0 represents False and 1 is True. |
| VideoAdapter.NumberOfColorPlanes | Display_Adapter_Number of Color Planes | ZENworks.VideoAdapter.NumberOfColorPlanes | Unsigned Integer | | Number of color planes supported by the video system. |
| VideoAdapter.CurrentVerticalResolution | Display_Adapter_Current Vertical Resolution | ZENworks.VideoAdapter.CurrentVertical Resolution | Unsigned Integer | | Vertical resolution of the display. |
| VideoAdapter.CurrentHorizontalResolution | Display_Adapter_Current Horizontal Resolution | ZENworks.VideoAdapter.CurrentHorizontal Resolution | Unsigned Integer | | Horizontal resolution of the display. |
| VideoAdapter.Description | Display_Adapter_Description | ZENworks.VideoAdapter.Description | String | 254 | Video adapter description. |

| Export Wizard Attribute Name | Export Attribute Name (Column Heading in the .csv file) | Database Schema Attribute Name | Data Type | Length | Description of the Attribute |
|----------------------------------|---|---|----------------------------------|--------|--|
| VideoAdapter.MinRefreshRate | Display Adapter_Minimum Refresh Rate | ZENworks.VideoAdapter.MinRefreshRate | Unsigned Integer | | Minimum refresh rate of the monitor for redrawing the display, measured in Hertz. |
| VideoAdapter.MaxRefreshRate | Display Adapter_Maximum Refresh Rate | ZENworks.VideoAdapter.MaxRefreshRate | Unsigned Integer | | Maximum refresh rate of the monitor for redrawing the display, measured in Hertz. |
| VideoAdapter.VideoArchitecture | Display Adapter_Video Architecture | ZENworks.VideoAdapter.VideoArchitecture | Unsigned Integer (enum) | | The architecture of the video subsystem in this system. For example, CGA/VGA/SVGA/8514A. See Section M.5, "Enumeration Values for Hardware-Display Adapter-Video Architecture," on page 769. |
| VideoAdapter.VideoMemoryType | Display Adapter_Video Memory Type | ZENworks.VideoAdapter.VideoMemoryType | Unsigned Small Integer (Enum) | | The type of memory for this adapter. For example, VRAM/SRAM/DRAM/EDO RAM. See Section M.6, "Enumeration Values for Hardware-Display Adapter-Video Memory Type," on page 769. |
| VideoAdapter.MaxMemorySupported | Display Adapter_Maximum Memory Supported(KB) | ZENworks.VideoAdapter.MaxMemorySupported | Unsigned Integer | | Maximum memory that the display adapter supports for VIDEO RAM. |
| VideoAdapter.CurrentBitsPerPixel | Display Adapter_Current Bits/Pixel | ZENworks.VideoAdapter.CurrentBitsPerPixel | Unsigned Integer | | Number of adjacent color bits for each pixel. |
| VideoAdapter.ChipSet | Display Adapter_Chip Set | ZENworks.VideoAdapter.ChipSet | String | 254 | The chip set used in the video adapter. |
| VideoAdapter.DACType | Display Adapter_DAC Type | ZENworks.VideoAdapter.DAC Type | String | 254 | The digital to analog converter type used in the video adapter. |
| VideoAdapter.ProviderName | Display Adapter_Provider | ZENworks.VideoAdapter.Provider | String | 254 | The manufacturer or the provider name. |
| ZENPOTSModem.Caption | Modem_Caption | ZENworks.ZENPOTSModem.Caption | String | 64 | The short name of the modem. |

| Export Wizard Attribute Name | Export Attribute Name (Column Heading in the .csv file) | Database Schema Attribute Name | Data Type | Length | Description of the Attribute |
|------------------------------|---|-----------------------------------|--|--------|---|
| ZENPOTSModem.Description | Modem_Description | ZENworks.ZENPOTSModem.Description | String | 254 | The complete description of the modem. For example, Standard 2400 bps modem, IBM PCMCIA HPC modem. |
| ZENPOTSModem.Name | Modem_Name | ZENworks.ZENPOTSModem.Name | String | 254 | The name of the modem dictating its type and usage. For example, Standard Windows Modem means that this is used in standard Windows architecture. |
| ZENPOTSModem.ProviderName | Modem_Provider | ZENworks.ZENPOTSModem.Provider | String | 254 | The manufacturer or the provider name. |
| ZENPOTSModem.DeviceID | Modem_DeviceID | ZENworks.ZENPOTSModem.DeviceID | String | 64 | The unique ID assigned to the device. |
| BIOS.BIOSIDBytes | BIOS_BIOS Identification Bytes | ZENworks.BIOS.BIOSIDBytes | String | 254 | Byte in the BIOS that indicates the computer model. |
| BIOS.SerialNumber | BIOS_Serial Number | ZENworks.BIOS.SerialNumber | String | 64 | Serial number of BIOS assigned by the manufacturer. |
| BIOS.PrimaryBIOS | BIOS_Primary Bios | ZENworks.BIOS.PrimaryBIOS | BIT (Used for Boolean conditions here) | | True when set to 1, indicating that this BIOS is the primary BIOS. Used in systems with additional BIOS chips. |
| BIOS.InstallDate | BIOS_Install Date | ZENworks.BIOS.InstallDate | String | 25 | The release date of the BIOS given by the manufacturer. |
| BIOS.Version | BIOS_Version | ZENworks.BIOS.Version | String | 254 | Version or revision level of the BIOS. |
| BIOS.Manufacturer | BIOS_Manufacturer | ZENworks.BIOS.Manufacturer | String | 254 | The manufacturer name of BIOS. |
| BIOS.Caption | BIOS_Caption | ZENworks.BIOS.Caption | String | 64 | The name of the BIOS as given by the BIOS manufacturer. |
| BIOS."size" | BIOS_Size(KB) | ZENworks.BIOS.size | Unsigned Integer | | Size of the BIOS in bytes. |

| Export Wizard Attribute Name | Export Attribute Name (Column Heading in the .csv file) | Database Schema Attribute Name | Data Type | Length | Description of the Attribute |
|----------------------------------|---|--------------------------------------|-------------------------------|--------|--|
| Processor.CurrentClockSpeed | Processor_Current Clock Speed(MHz) | CIM.Processor.CurrentClockSpeed | Unsigned Integer | | Current clock speed of the processor in MHz. |
| Processor.MaxClockspeed | Processor_Maximum Clock Speed(MHz) | CIM.Processor.MaxClockSpeed | Unsigned Integer | | Maximum clock speed of the processor in MHz. |
| Processor.Role | Processor_Role | CIM.Processor.Role | String | 254 | Type of processor such as central processor, math coprocessor, and others |
| Processor.Family | Processor_Family | CIM.Processor.Family | Unsigned Small Integer (enum) | | Family the processor belongs to. See Section M.9, "Enumeration Values for Hardware-Processor-Family," on page 770. |
| Processor.Otherfamilydescription | Processor_Other Family Description | CIM.Processor.OtherFamilyDescription | String | 64 | Additional description about the processor family, such as the Pentium processor with MMX technology when the processor cannot be designated using Family. |
| Processor.UpgradeMethod | Processor_Upgrade Method | CIM.Processor.UpgradeMethod | Unsigned Small Integer (Enum) | | The method by which this processor can be upgraded, if upgrades are supported. See Section M.10, "Enumeration Values for Hardware-Processor-Upgrade Method," on page 770. |
| Processor.Stepping | Processor_Processor Stepping | CIM.Processor.Stepping | String | 254 | Single-byte code characteristic provided by microprocessor vendors to identify the processor stepping model. |
| Processor.DeviceID | Processor_DeviceID | CIM.Processor.DeviceID | String | 64 | Special hexadecimal string identifying the processor type. |
| CacheMemory.Speed | Cache Memory_Speed(nsec) | CIM.PhysicalMemory.Speed | Unsigned Integer | | Speed of this System Cache module in nanoseconds. This is stored in CIM.PhysicalMemory class and is associated to CIM.CacheMemory. For more information on how they are associated, see "Understanding the ZENworks 7 Server Managements Inventory Database Schema" on page 549. |

| Export Wizard Attribute Name | Export Attribute Name (Column Heading in the .csv file) | Database Schema Attribute Name | Data Type | Length | Description of the Attribute |
|-------------------------------|---|-----------------------------------|----------------------------------|--------|---|
| CacheMemory.Capacity | Cache Memory_Capacity(MB) | CIM.PhysicalMemory.Capacity | Unsigned Integer | | Capacity of this System Cache module in nanoseconds. This is stored in CIM.PhysicalMemory class and is associated to CIM.CacheMemory. For more information on how they are associated, see "Understanding the ZENworks 7 Server Managements Inventory Database Schema" on page 549. |
| CacheMemory.Level | Cache Memory_Level | CIM.CacheMemory."Level" | Unsigned Small Integer (enum) | | Indicates the cache level: internal cache that is built in to the microprocessors, or external cache that is between the CPU and DRAM. |
| CacheMemory.WritePolicy | Cache Memory_Write Policy | CIM.CacheMemory.WritePolicy | Unsigned Small Integer (enum) | | Indicates the two different ways (Write-Back and Write-Through Cache) that the cache can handle to write to the memory. |
| CacheMemory.Errormethodology | Cache Memory_Error Methodology | CIM.CacheMemory.Error Methodology | String | 254 | Error correction scheme supported by this cache component, for example, Parity/Single Bit ECC/MultiBit ECC. |
| CacheMemory.CacheType | Cache Memory_Cache Type | CIM.CacheType | Unsigned Small Integer (enum) | | Defines the system cache type. For example, Instruction, Data, Unified. |
| CacheMemory.LineSize | Cache Memory_Line Size (Bytes) | CIM.CacheMemory.LineSize | Unsigned Integer | | Size in bytes of a single cache bucket or line. |
| CacheMemory.ReplacementPolicy | Cache Memory_Replacement Policy | CIM.CacheMemory.ReplacementPolicy | Unsigned Integer (enum) | | Algorithm that the cache uses to determine which cache lines or buckets should be reused. |
| CacheMemory.ReadPolicy | Cache Memory_Read Policy | CIM.CacheMemory.ReadPolicy | Unsigned Small Integer (enum) | | Indicates whether the data cache is for read operation. |
| CacheMemory.Associativity | Cache Memory_Associativity | CIM.CacheMemory.Associativity | Unsigned Integer (enum) | | Defines the system cache associativity (direct-mapped, 2-way, 4-way). |

| Export Wizard Attribute Name | Export Attribute Name (Column Heading in the .csv file) | Database Schema Attribute Name | Data Type | Length | Description of the Attribute |
|----------------------------------|---|---|------------------------|--------|--|
| Diskette Drive.Manufacturer | Diskette Drive_Manufacturer | ZENworks.PhysicalDiskette.Manufacturer | String | 254 | Vendor name. |
| Diskette Drive.Description | Diskette Drive_Description | ZENworks.PhysicalDiskette.Description | String | 254 | Floppy diskette description. |
| Diskette Drive.PhysicalCylinders | Diskette Drive_PhysicalCylinders | ZENworks.PhysicalDiskette.PhysicalCylinders | Unsigned Integer | | Total number of cylinders or tracks on the floppy. |
| Diskette Drive.PhysicalHeads | Diskette Drive_PhysicalHeads | ZENworks.PhysicalDiskette.PhysicalHeads | Unsigned Small Integer | | Number of heads. |
| Diskette Drive.Capacity | Diskette Drive_Capacity (MB) | ZENworks.PhysicalDiskette.Capacity | Unsigned Integer | | Total size. |
| Diskette Drive.SectorsPerTrack | Diskette Drive_Sectors/Track | ZENworks.PhysicalDiskette.SectorsPerTrack | Unsigned Integer | | Number of sectors per track. |
| Diskette Drive.DeviceID | Diskette Drive_DeviceID | CIM.Diskette Drive | String | 64 | The drive name representing the floppy drive. |
| ZENDiskDrive.Manufacturer | Physical Disk Drive_Manufacturer | ZENworks.PhysicalDisk.Manufacturer | String | 254 | Vendor name. |
| ZENDiskDrive.Description | Physical Disk Drive_Description | ZENworks.PhysicalDisk.Description | String | 254 | Hard disk vendor description. |
| ZENDiskDrive.PhysicalCylinders | Physical Disk Drive_PhysicalCylinders | ZENworks.PhysicalDisk.PhysicalCylinders | Unsigned Integer | | Total number of cylinders. |
| ZENDiskDrive.PhysicalHeads | Physical Disk Drive_PhysicalHeads | ZENworks.PhysicalDisk.PhysicalHeads | Unsigned Small Integer | | Number of heads. |

| Export Wizard Attribute Name | Export Attribute Name (Column Heading in the .csv file) | Database Schema Attribute Name | Data Type | Length | Description of the Attribute |
|------------------------------------|---|--|------------------|--------|---|
| ZENDiskDrive.SectorsPerTrack | Physical Disk Drive_Sectors/Track | ZENworks.PhysicalDisk.SectorsPerTrack | Unsigned Integer | | Number of sectors per track. |
| ZENDiskDrive.Capacity | Physical Disk Drive_Capacity(MB) | ZENworks.PhysicalDisk.Capacity | Unsigned Integer | | Total size of the hard disk. |
| ZENDiskDrive.Removable | Physical Disk Drive_Removable | ZENworks.LogicalDiskDrive.Removable | BIT | | 0 indicates that it is a fixed disk and 1 indicates that it is a removable disk. |
| LocalFileSystem.DeviceID | Logical Disk Drive_Device ID | ZENworks.LogicalDiskDrive.DeviceID | String | 64 | The drive letter, such as C: or A:. |
| LocalFileSystem.FileSystemSize | Logical Disk Drive_Size(MB) | CIM.LocalFileSystem.FileSystemSize | Integer | | The total size of the file system or the logical disk. |
| LocalFileSystem.AvailableSpace | Logical Disk Drive_Free Size(MB) | CIM.LocalFileSystem.AvailableSpace | Integer | | The available size of the file system or the logical disk. |
| LocalFileSystem.VolumeSerialNumber | Logical Disk Drive_Volume Serial Number | CIM.LocalFileSystem.VolumeSerialNumber | String | 254 | The volume serial number of the specified drive. |
| LocalFileSystem.Caption | Logical Disk Drive_Caption | CIM.LocalFileSystem.Caption | String | 64 | The volume label of the specified drive. |
| LocalFileSystem.FileSystemType | Logical Disk Drive_File System Type | CIM.LocalFileSystem.FileSystemType | String | 254 | The file system on the drive, such as FAT or NTFS. |
| CDROMDrive.Manufacturer | CDROM_Manufacturer | ZENworks.PhysicalCDROM.Manufacturer | String | 254 | The manufacturer of the CD-ROM drive. |
| CDROMDrive.Caption | CDROM_Caption | ZENworks.PhysicalCDROM.Caption | String | 64 | CD-ROM label. |
| CDROMDrive.Description | CDROM_Description | ZENworks.PhysicalCDROM.Description | String | 254 | Description of the CD drive, as given by the manufacturer. For example, ATAPI CDROM, CREATIVE CD1620E SL970520. |

| Export Wizard Attribute Name | Export Attribute Name (Column Heading in the .csv file) | Database Schema Attribute Name | Data Type | Length | Description of the Attribute |
|------------------------------|---|-----------------------------------|--|--------|---|
| CDROMDrive.DeviceID | CDROM_Device ID | ZENworks.Logical CDROM.DeviceID | String | 64 | Drive letter allocated for the CD on the inventoried server. |
| SerialPort.Name | Serial Port_Name | ZENworks.SerialPort.Name | String | 254 | The name of the serial port. For example, COM1, COM2, and others. |
| SerialPort.Address | Serial Port_Address | ZENworks.SerialPort.Address | Unsigned Integer | | The address mapped in memory for the serial port. |
| SerialPort.IRQNumber | Serial Port_IRQ Number | CIM.IRQ.IRQNumber | Unsigned Integer | | The IRQ channel on the system to which the serial port is attached. In the database, this information is stored in an IRQ class and not in a Serial Port class. For more information on how they are associated, see Chapter 15, "Understanding the ZENworks 7 Server Managements Inventory Database Schema," on page 549. |
| ParallelPort.Name | Parallel Port_Name | ZENworks.ParallelPort.Name | String | 254 | The name of the parallel port. For example, LPT1 and others. |
| ParallelPort.Address | Parallel Port_Address | ZENworks.ParallelPort.Address | Unsigned Integer | | The name of the parallel port. For example, LPT1 and others. |
| ParallelPort.DMA Support | Parallel Port_DMA Support | ZENworks.ParallelPort.DMA Support | BIT (used for Boolean conditions here) | | If True or 1, then it means that DMA is the channel that is allocated for bulk data transfer for use with devices connected to the parallel ports. |

| Export Wizard Attribute Name | Export Attribute Name (Column Heading in the .csv file) | Database Schema Attribute Name | Data Type | Length | Description of the Attribute |
|------------------------------------|---|--|------------------|--------|--|
| ParallelPort.IRQNumber | ParallelPort_IRQNumber | CIM.IRQ.IRQNumber | Unsigned Integer | | The IRQ channel on the system to which the parallel port is attached. This information is stored in an IRQ class and not in a parallel port class in the database. For more information on how they are associated, see Chapter 15, "Understanding the ZENworks 7 Server Managements Inventory Database Schema," on page 549. |
| Bus.Version | Bus_Version | ZENworks.Bus.BusVersion | String | 254 | Version of the bus supported by the inventoried server. |
| Bus.Description | Bus_Description | ZENworks.Bus.Description | String | 254 | Description of the bus. |
| Bus.BusType | Bus_Bus Type | ZENworks.Bus.BusType | Integer (enum) | | The bus type of the system. |
| Bus.Name | Bus_Name | ZENworks.Bus.Name | String | 254 | Name of the internal system bus. |
| Bus.DeviceID | Bus_Device ID | ZENworks.Bus.DeviceID | String | 64 | The unique ID for the specific bus. |
| ZENNetworkAdapter.Name | NetworkAdapter_Name | CIM.ZENworks.ZENAdapter.Name | String | 254 | Network adapters installed on the system. |
| ZENNetworkAdapter.MaxSpeed | NetworkAdapter_Max_Speed (Mbps) | CIM.ZENworks.ZENAdapter.MaxSpeed | Unsigned Integer | | Rate at which the adapter can transfer data. |
| ZENNetworkAdapter.PermanentAddress | NetworkAdapter_Permanent Address | CIM.ZENworks.ZENAdapter.PermanentAddress | String | 64 | Machine address stored permanently in the adapter (MAC address). |
| ZENNetworkAdapter.MACAddress | NetworkAdapter_Address | CIM.ZENworks.ZENAdapter.MACAddress | String | 64 | The MAC address stored in the network adapter. |
| ZENNetworkAdapter.ProviderName | NetworkAdapter_Provider | CIM.ZENworks.ZENAdapter.Provider | String | 254 | The manufacturer or the provider. |
| ZENNetworkAdapter.AdapterType | NetworkAdapter_Adapter Type | CIM.ZENworks.ZENAdapter.AdapterType | String | 254 | Type of the adapter, such as Ethernet or FDDI adapter. |

| Export Wizard Attribute Name | Export Attribute Name (Column Heading in the .csv file) | Database Schema Attribute Name | Data Type | Length | Description of the Attribute |
|------------------------------|---|------------------------------------|------------------------|--------|--|
| SoundAdapter.Description | Multimedia Card_Description | ZENworks.SoundAdapter.Description | String | 254 | Description of the multimedia component for the inventoried server. |
| SoundAdapter.Name | Multimedia Card_Name | ZENworks.SoundAdapter.Name | String | 254 | Name of the sound card installed on the system. |
| SoundAdapter.Manufacturer | Multimedia Card_Manufacturer | ZENworks.SoundAdapter.Manufacturer | String | 254 | Vendor name. |
| SoundAdapter.ProviderName | Multimedia Card_Provider | ZENworks.SoundAdapter.Provider | String | 254 | The provider or the manufacturer of the multimedia card. |
| Battery.Name | Battery_Name | CIM.Battery.Name | String | 254 | Name of the battery installed on the system. |
| Battery.Chemistry | Battery_Chemistry | CIM.Battery.Chemistry | Unsigned Small Integer | | Indicates the battery's chemistry, such as lead acid, nickel cadmium and others. See Section M.8, "Enumeration Values for Hardware-Battery-Chemistry," on page 769. |
| Battery.DesignCapacity | Battery_DesignCapacity(mWatt-hours) | CIM.Battery.DesignCapacity | Unsigned Integer | | The design capacity of the battery in mWatt-hours. |
| Battery.DesignVoltage | Battery_DesignVoltage(Millivolts) | CIM.Battery.DesignVoltage | Unsigned Integer | | The design voltage of the battery in mVolts. |
| Battery.SmartBatteryVersion | Battery_Smart BatteryVersion | CIM.Battery.SmartBatteryVersion | String | 64 | The Smart Battery Data Specification version number supported by this battery. |
| Battery.Manufacturer | Battery_Manufacturer | CIM.PhysicalComponent.Manufacturer | String | 254 | Vendor name of the battery. |
| Battery.InstallDate | Battery_InstallDate | CIM.PhysicalComponent.InstallDate | String | 25 | Date of manufacturing the battery. |
| Battery.SerialNumber | Battery_SerialNumber | CIM.PhysicalComponent.SerialNumber | String | 64 | Battery serial number. |
| PowerSupply.Description | Power Supply_Description | CIM.PowerSupply.Description | String | 254 | Name and description of the power supply on the system. |

| Export Wizard Attribute Name | Export Attribute Name (Column Heading in the .csv file) | Database Schema Attribute Name | Data Type | Length | Description of the Attribute |
|-------------------------------|---|-------------------------------------|------------------|--------|---|
| PowerSupply.TotalOutputPower | Power Supply_Total Output Power (MilliWatts) | CIM.Power Supply.Total OutputPower | Unsigned Integer | | Total output power of the power supply. |
| IPProtocolEndpoint.Address | IP Address_Address | CIM.IP Protocol Endpoint.Address | String | 254 | IP address of the inventoried server. |
| IPProtocolEndpoint.SubnetMask | IP Address_Subnet Mask | CIM.IP Protocol Endpoint.SubnetMask | String | 254 | The subnet mask of the inventoried server. |
| DNSName.LABEL | DNS_LABEL | ManageWise.DNSName.Label | String | 254 | DNS name of the inventoried server. |
| IPXProtocolEndpoint.Address | IPX Address_Address | CIM.IPX Protocol Endpoint.Address | String | 254 | IPX address of the inventoried server. |
| LANEndPoint.MACAddress | MAC Address_Address | CIM.LAN Endpoint.MACAddress | String | 12 | MAC address of the inventoried server. |
| MotherBoard.Version | MotherBoard_Version | ZENworks.Motherboard.Version | String | 64 | Motherboard version. |
| MotherBoard.Description | MotherBoard_Description | ZENworks.Motherboard.Description | String | 254 | The description of the motherboard. |
| MotherBoard.Manufacturer | MotherBoard_Manufacturer | ZENworks.Motherboard.Manufacturer | String | 254 | The manufacturer of the motherboard. |
| MotherBoard.NumberOfSlots | MotherBoard_Number Of Slots | ZENworks.Motherboard.Numberofslots | Integer | | The number of expansion slots on the motherboard. |
| IRQ.Number | IRQ_IRQ Number | CIM.IRQ.IRQN umber | Unsigned Integer | | The system interrupt number. |

| Export Wizard Attribute Name | Export Attribute Name (Column Heading in the .csv file) | Database Schema Attribute Name | Data Type | Length | Description of the Attribute |
|------------------------------|---|--------------------------------|----------------------------------|--------|---|
| IRQ.Availability | IRQ_Availability | CIM.IRQ.Availability | Unsigned Small Integer (Enum) | | Indicates whether the IRQ channel is used or available. Enumeration values are as follows: 1 = "Other" 2 = "Unknown" 3 = "Available" 4 = "In Use/Not Available" 5 = "In Use and Available/Shareable" |
| IRQ.TriggerType | IRQ_IRQ Trigger Type | CIM.IRQ.TriggerType | Unsigned Small Integer | | IRQ trigger type indicating whether edge (value=4) or level triggered (value=3) interrupts occur. Enumeration values are as follows: 1 = "Other" 2 = "Unknown" 3 = "Level" 4 = "Edge" |
| IRQ.Shareable | IRQ_IRQ Shareable | CIM.IRQ.Shareable | Unsigned Small Integer | | Boolean indicating whether the IRQ can be shared. |
| SLOT.MaxData Width | Slot_Maximum Data Width | CIM.Slot.MaxData Width | Unsigned Small Integer | | Maximum bus width of adapter cards that can be inserted into this slot in bits. If the value is Unknown, enter 0. If the value is other than 8, 16, 32, 64 or 128, enter 1. It is expressed in bits. |
| SLOT.ThermalRating | Slot_Thermal Rating (MilliWatts) | CIM.Slot.Thermal Rating | Unsigned Integer | | Maximum thermal dissipation of the slot in milliwatts. |
| SLOT.Description | Slot_Description | CIM.SlotDescription | String | 254 | The description of the adapter mounted on the slot. |
| DMA.DMAChannel | DMA_DMA Channel Number | CIM.DMA.DMAChannel | Unsigned Integer | | The DMA channel number. |
| DMA.Description | DMA_Description | CIM.DMA.Description | String | 254 | The name of the device using the DMA channel. |

| Export Wizard Attribute Name | Export Attribute Name (Column Heading in the .csv file) | Database Schema Attribute Name | Data Type | Length | Description of the Attribute |
|--------------------------------|---|--|--|--------|--|
| DMA.Availability | DMA_Availability | CIM.DMA.Availability | Unsigned Small Integer | | Indicates whether the DMA channel is available. Enumeration values are as follows: 1 = "Other" 2 = "Unknown" 3 = "Available" 4 = "In Use/Not Available" 5 = "In Use and Available/Shareable" |
| DMA.BurstMode | DMA_DMA Burst Mode | CIM.DMA.BurstMode | BIT (used for Boolean condition here) | | Indication that the DMA channel supports the burst mode. |
| NetWareOperatingSystem.Version | NetWare.Version | ZENworks.NetWareOperating.Version | String | 254 | Version of the NetWare operating system. |
| Memory.TotalMemory | Memory_TotalMemory(MB) | ZENOperatingSystem.TotalVisibleMemorySize | Integer | | Total memory of the Windows operating system. |
| MSDomainName.Label | WindowsDomain_Name | ManageWise.MSDomainName | String | 254 | The Windows domain to which the server is attached. |
| Monitor.DeviceID | Monitor_DeviceID | ZENworks.ZENDesktopMonitor.DeviceID | Integer | | Unique ID of a desktop monitor that is attached to a computer system. |
| Monitor.Description | Monitor_Description | ZENworks.ZENDesktopMonitor.Description | varchar | 254 | Description of the monitor. |
| Monitor.ModelID | Monitor_ModelID | ZENworks.ZENDesktopMonitor.ModelID | varchar | | Unique ID of a model of the monitor. It is a combination of the Manufacturer ID and Product ID. |
| Monitor.ManufactureDate | Monitor_ManufactureDate | ZENworks.ZENDesktopMonitor.ManufactureDate | char | 25 | Year in which the monitor was manufactured. |
| Monitor.ViewableSize | Monitor_ViewableSize | ZENworks.ZENDesktopMonitor.ViewableSize | integar | | A number representing the diagonal width of the screen image excluding the black borders around the image's edge. |

| Export Wizard Attribute Name | Export Attribute Name (Column Heading in the .csv file) | Database Schema Attribute Name | Data Type | Length | Description of the Attribute |
|------------------------------|---|--|--------------------|--------|--|
| Monitor.NominalSize | Monitor_NominalSize | ZENworks.ZENDesktopMonitor.NominalSize | integer | | A number representing the diagonal width of the monitor (the distance from one corner of the screen to the opposite corner of the screen). |
| Monitor.SerialNumber | Monitor_SerialNumber | ZENworks.ZENDesktopMonitor.SerialNumber | varchar | 128 | Manufacturer's number used to identify a monitor. |
| Monitor.Manufacturer | Monitor_Manufacturer | ZENworks.ZENDesktopMonitor.Manufacturer | varchar | 254 | Name of the monitor's manufacturer. |
| Monitor.Model | Monitor_Model | ZENworks.ZENDesktopMonitor.Model | varchar | 254 | Product name of the monitor given by the manufacturer. |
| Chassis.AssetTag | Chassis_AssetTag | ZENworks_ZENChassis | varchar | 254 | Asset tag number of the system chassis. |
| Chassis.ChassisType | Chassis_ChassisType | ZENworks_ZENChassis | unsigned small int | | Represents whether the system chassis is a laptop, desktop, notebook, docking station and so on. |
| Chassis.NumberOfPowerCords | Chassis_NumberOfPowerCords | ZENworks_ZENChassis | varchar | 128 | Total number of power cords attached to a system chassis. |
| Chassis.Manufacturer | Chassis_Manufacturer | ZENworks_ZENChassis | varchar | 254 | Name of the system chassis manufacturer. |
| Chassis.SerialNumber | Chassis_SerialNumber | ZENworks_ZENChassis | varchar | 128 | Manufacturer's number used to identify a system chassis. |
| Chassis.Version | Chassis_Version | ZENworks_ZENChassis | varchar | 64 | Version number of the system chassis. |
| Chassis.Tag | Chassis_Tag | ZENworks_ZENChassis | varchar | 64 | Unique ID of the system chassis attached to a particular computer system. |
| Software.ProductIdentifier | Software_productIdentifier | MW_DBA.InstalledSoftware.productIdentifier | varchar | 254 | A unique, 16-character identifier for an installed product. This identifier is available from MSI on Windows. |
| Software.InternalVersion | Software_InternalVersion | MW_DBA.InstalledSoftware.InternalVersion | varchar | 64 | Internal version of a product |
| Software.Language | Software_Language | MW_DBA.InstalledSoftware.Language | smallint | | User-friendly name for the language of this copy of the product. |

| Export Wizard Attribute Name | Export Attribute Name (Column Heading in the .csv file) | Database Schema Attribute Name | Data Type | Length | Description of the Attribute |
|------------------------------|---|--|-----------|--------|---|
| Software.UninstallString | Software_UninstallString | MW_DBA.InstalledSoftware.UninstallString | varchar | 254 | The command to invoke for uninstalling this product instance. Currently, this is available in Add/Remove Programs (ARP) and MSI on Windows. |
| Software.InstallationSource | Software_InstallationSource | MW_DBA.InstalledSoftware.InstallationSource | varchar | 254 | Identifies the file system path where the installation files were stored when installing this product instance. Currently, this is available in ARP and MSI on Windows. |
| Software.FriendlyName | Software_FriendlyName | MW_DBA.InstalledSoftware.FriendlyName | varchar | 254 | Display name of the software. |
| Software.LastExecutionTime | Software_LastExecutionTime | MW_DBA.InstalledSoftware.LastExecutionTime | bigint | | Date and time stamp when the product was last executed. |
| Software.FrequencyOfUsage | Software_FrequencyOfUsage | MW_DBA.InstalledSoftware.FrequencyOfUsage | smallint | | Number of times the product is used. |
| Software.Description | Software_Description | MW_DBA.InstalledSoftware.Description | varchar | 254 | Description of the product. |
| Software.DefinitionDate | Software_DefinitionDate | MW_DBA.InstalledVirusScanner.DefinitionDate | bigint | | The date of the virus definition file installed on the computer. Some anti-virus products combine date and version into a single string. |
| Software.DefinitionVersion | Software_DefinitionVersion | MW_DBA.InstalledVirusScanner.DefinitionVersion | varchar | 64 | The vendor-defined version of the virus definition file that has been installed on a computer |
| Software.Edition | Software_Edition | MW_DBA.ProductEdition.Name | varchar | 128 | Product edition defined by the vendor. For example, Professional. |
| Software.SupportPack | Software_SupportPack | MW_DBA.SupportPack.Name | varchar | 128 | Support pack name. |
| Software.Path | Software_Path | MW_DBA.Directory.Path | varchar | 254 | Directory path where the product is installed on the computer system. |
| Software.Name | Software_Name | MW_DBA.Software.Name | varchar | 254 | Vendor-defined name of the product represented as a vendor trademark or registered trademark. |

| Export Wizard Attribute Name | Export Attribute Name (Column Heading in the .csv file) | Database Schema Attribute Name | Data Type | Length | Description of the Attribute |
|------------------------------|---|----------------------------------|-----------|--------|--|
| Software.Vendor | Software_Vendor | MW_DBA.Software.Vendor | varchar | 254 | Name of the software manufacturer |
| Software.Version | Software_Version | MW_DBA.Software.Version | varchar | 64 | User-friendly version of a product. |
| Software.Category | Software_Category | MW_DBA.Software.Category | varchar | 64 | Product category to which the product belongs. |
| Software.HelpLink | Software_HelpLink | MW_DBA.Software.HelpLink | varchar | 254 | Support web site URL for the product that is available in ARP and MSI. |
| Software.PackageGUID | Software_PackageGUID | MW_DBA.Software.PackageGUID | varchar | 64 | Vendor-defined GUID for a product that is available in MSI. |
| Software.Patch Name | Software_Patch Name | MW_DBA.Patch.Name | varchar | 254 | Vendor-defined name for the patch. |
| File.Name | File_Name | MW_DBA.File.Name | varchar | 254 | Name of the file representing the software. |
| File.FileVersion | File_FileVersion | MW_DBA.File.FileVersion | varchar | 64 | Version of the file representing the software. |
| File.InternalName | File_InternalName | MW_DBA.File.InternalName | varchar | 254 | Internal name. |
| File.ProductVersion | File_ProductVersion | MW_DBA.File.ProductVersion | varchar | 64 | The version of the product represented by this file. |
| File.Size | File_size | MW_DBA.File.size | bigint | | Size of the file representing the software. |
| File.LastModified | File_LastModified | MW_DBA.File.LastModified | bigint | | Last modified date of the file representing the software. |
| File.Company | File_Company | MW_DBA.File.Company | varchar | 254 | Vendor name. |
| File.ProductName | File_ProductName | MW_DBA.File.ProductName | varchar | 254 | The product which this file represents. |
| File.Language | File_Language | MW_DBA.File.Language | smallint | | User-friendly name for the language of this copy of the file |
| File.SoftwareDictionaryID | File_SoftwareDictionaryID | MW_DBA.File.SoftwareDictionaryID | varchar | 64 | ID of the file as represented in the General software dictionary. |
| DiskUsage.TotalDiskUsage | DiskUsage.TotalDiskUsage | MW_DBA.DiskUsage.TotalDiskUsage | bigint | | Total disk usage for all the files of the specified extension. |
| DiskUsage.Name | DiskUsage.Name | MW_DBA.DiskUsage.Name | varchar | 32 | The file extension for which the disk usage is scanned for. |

Enumeration Values

M

This section provides information on the following topics:

- ◆ [Section M.1, “Enumeration Values for General-System Information-Management Technology,” on page 767](#)
- ◆ [Section M.2, “Enumeration Values for General-Inventory Information-Scan Mode,” on page 768](#)
- ◆ [Section M.3, “Enumeration Values for Software-Operating Systems-Windows - Name,” on page 768](#)
- ◆ [Section M.4, “Enumeration Values for Installation Repository,” on page 768](#)
- ◆ [Section M.5, “Enumeration Values for Hardware-Display Adapter-Video Architecture,” on page 769](#)
- ◆ [Section M.6, “Enumeration Values for Hardware-Display Adapter-Video Memory Type,” on page 769](#)
- ◆ [Section M.7, “Enumeration Values for Hardware-Pointing Device-Name,” on page 769](#)
- ◆ [Section M.8, “Enumeration Values for Hardware-Battery-Chemistry,” on page 769](#)
- ◆ [Section M.9, “Enumeration Values for Hardware-Processor-Processor Family,” on page 770](#)
- ◆ [Section M.10, “Enumeration Values for Hardware-Processor-Upgrade Method,” on page 770](#)
- ◆ [Section M.11, “Enumeration Values for Hardware-Chassis-Chassis Type,” on page 770](#)
- ◆ [Section M.12, “Enumeration Values for Hardware-Bus-Protocol Supported,” on page 771](#)
- ◆ [Section M.13, “Enumeration Values for Hardware-Processor-Role,” on page 771](#)
- ◆ [Section M.14, “Enumeration Values for System-System Cache-Level,” on page 771](#)
- ◆ [Section M.15, “Enumeration Values for System-System Cache-Cache Type,” on page 771](#)
- ◆ [Section M.16, “Enumeration Values for System-System Cache-Replacement Policy,” on page 771](#)
- ◆ [Section M.17, “Enumeration Values for System-System Cache-Read Policy,” on page 772](#)
- ◆ [Section M.18, “Enumeration Values for System-System Cache-Write Policy,” on page 772](#)
- ◆ [Section M.19, “Enumeration Values for System-System Cache-Associativity,” on page 772](#)
- ◆ [Section M.20, “Enumeration Values for System-System IRQ-Availability,” on page 772](#)
- ◆ [Section M.21, “Enumeration Values for System-System IRQ-IRQ Trigger Type,” on page 772](#)
- ◆ [Section M.22, “Enumeration Values for System-System DMA-Availability,” on page 773](#)
- ◆ [Section M.23, “Enumeration Values for Language,” on page 773](#)

M.1 Enumeration Values for General-System Information-Management Technology

1 = Unknown

3 = DMI Enabled

5 = SNMP Enabled

| | | |
|-----------|-----------------|-------------------------|
| 2 = Other | 4 = WMI Enabled | 6 = DMI and WMI Enabled |
|-----------|-----------------|-------------------------|

M.2 Enumeration Values for General-Inventory Information-Scan Mode

| | | |
|-------------|---------|-----------------|
| 1 = Unknown | 3= DMI | 5= SNMP |
| 2 = Other | 4 = WMI | 6 = DMI and WMI |

M.3 Enumeration Values for Software-Operating Systems-Windows - Name

| | | |
|-------------|--------------|---------------------|
| 0 = Unknown | 18 = WINNT | 59 = Dedicated |
| 1 = Other | 21 = NetWare | 63 = Windows (R) Me |
| 16 = WIN95 | 36 = Linux | 67 = Windows XP |
| 17 = WIN98 | 58 = Windows | |

M.4 Enumeration Values for Installation Repository

The following Installation Repository enum values are displayed in the Software > Software Group Components > Software Group and Software > Software Components > Software classes.

| | | |
|--|---|---|
| 0 = Others | 7 = MSI, Add Remove Programs, Software Dictionary | 20 = Software Dictionary, Probe |
| 1 = MSI | 8 = NetWare Products.dat | 21= MSI, Software Dictionary, Probes |
| 2 = Add/Remove Programs | 12 = Software Dictionary, NetWare Products.dat | 22 = Add Remove programs, Software Dictionary, Probe |
| 3 = MSI, Add Remove Programs | 16 = Probe | 23 = MSI, Add Remove programs, Software Dictionary, Probe |
| 4 = Software Dictionary | 17 = MSI, Probe | 24 = NetWare Products.dat, Probe |
| 5 = MSI, Software Dictionary | 18 = Add Remove Programs, Probe | 28 = Software Dictionary, NetWare Products.dat, Probe |
| 6 = Add Remove Programs, Software Dictionary | 19 = MSI, Add Remove Programs, Probe | |

M.5 Enumeration Values for Hardware-Display Adapter-Video Architecture

| | | |
|-------------|------------|--------------------------|
| 1 = Other | 6 = SVGA | 11 = XGA |
| 2 = Unknown | 7 = MDA | 12 = Linear Frame Buffer |
| 3 = CGA | 8 = HGC | 160 = PC-98 |
| 4 = EGA | 9 = MCGA | |
| 5 = VGA | 10 = 8514A | |

M.6 Enumeration Values for Hardware-Display Adapter-Video Memory Type

| | | |
|-------------|----------------------------|------------|
| 1 = Other | 6 = WRAM | 11 = 3DRAM |
| 2 = Unknown | 7 = EDO RAM | 12 = SDRAM |
| 3 = VRAM | 8 = Burst Synchronous DRAM | 13 = SGRAM |
| 4 = DRAM | 9 = Pipelined Burst SRAM | |
| 5 = SRAM | 10 = CDRAM | |

M.7 Enumeration Values for Hardware-Pointing Device-Name

| | | |
|-------------|-----------------|----------------------------|
| 1 = Other | 4 = Track Ball | 7 = Touch Pad |
| 2 = Unknown | 5 = Track Point | 8 = Touch Screen |
| 3 = Mouse | 6 = Glide Point | 9 = Mouse - Optical Sensor |

M.8 Enumeration Values for Hardware-Battery-Chemistry

| | |
|--------------------|--------------------------|
| 1 = Other | 5 = Nickel Metal Hydride |
| 2 = Unknown | 6 = Lithium-ion |
| 3 = Lead Acid | 7 = Zinc air |
| 4 = Nickel Cadmium | 8 = Lithium Polymer |

M.9 Enumeration Values for Hardware-Processor-Processor Family

| | | |
|---|---------------------------------------|--|
| 1 = Other | 24 = AMD Duron(TM) Processor Family | 130 = Itanium(TM) Processor |
| 2 = Unknown | 25 = K5 Family | 176 = Pentium(R) III Xeon(TM) |
| 11 = Pentium(R) Brand | 26 = K6 Family | 177= Pentium(R) III Processor with Intel(R) SpeedStep(TM) Technology |
| 12 = Pentium(R) Pro | 27 = K6 -2 | 178 = Pentium(R) 4 Processor |
| 13 = Pentium(R) II | 28 = K6 -3 | 181 = Inter(R) Xeon (TM) Processor MP |
| 14 = Pentium(R) Processor with MMX(TM) Technology | 29 = AMD Athlon (TM) Processor Family | 182 = AMD Athlon XP (TM) Processor Family |
| 15 = Celeron(TM) | 30 = AMD29000 Family | 183 = AMD Athlon MP(TM) Processor Family |
| 16 = Pentium(R) II Xeon(TM) | 31 = K6-2+ | 300 = 6 x 86 |
| 17 = Pentium(R) II | | |

M.10 Enumeration Values for Hardware-Processor-Upgrade Method

| | | |
|--------------------|----------------------------|---------------------|
| 1= Other | 5 = Replacement/Piggy Back | 9 = Slot 2 |
| 2 = Unknown | 6 = None | 10 = 370 Pin Socket |
| 3 = Daughter Board | 7 = LIF Socket | 11 = Slot A |
| 4 = ZIF Socket | 8 = Slot 1 | 12 = Slot M |

M.11 Enumeration Values for Hardware-Chassis-Chassis Type

| | | |
|-------------------------|----------------------|----------------------------|
| 1 = Other | 10 = Notebook | 19 = SubChassis |
| 2 = Unknown | 11 = Hand Held | 20 = Bus Expansion Chassis |
| 3 = Desktop | 12 = Docking Station | 21 = Peripheral Chassis |
| 4 = Low Profile Desktop | 13 = All in One | 22 = Storage Chassis |
| 5 = Pizza Box | 14 = Sub Notebook | 23 = Rack Mount Chassis |
| 6 = Mini Tower | 15 = Space-Saving | 24 = Sealed-Case PC |
| 7= Tower | 16 = Lunch Box | 25 = Multi-system Chassis |

| | |
|--------------|--------------------------|
| 8 = Portable | 17 = Main System Chassis |
| 9 = LapTop | 18 = Expansion Chassis |

M.12 Enumeration Values for Hardware-Bus-Protocol Supported

| | | |
|------------------|----------------|-----------------------------|
| 0 = Internal | 6 = VME Bus | 12 = Internal Processor |
| 1 = ISA | 7 = NuBus | 13 = Internal Power Bus |
| 2 = EISA | 8 = PCMCIA Bus | 14 = PNP ISA Bus |
| 3 = MicroChannel | 9 = C Bus | 15 = PNP Bus |
| 4 = TurboChannel | 10 = MPI Bus | 16 = Maximum Interface Type |
| 5 = PCI Bus | 11 = MPSA Bus | |

M.13 Enumeration Values for Hardware-Processor-Role

| | | |
|-------------|-----------------------|---------------------|
| 1 = Other | 3 = Central Processor | 5 = DSP Processor |
| 2 = Unknown | 4 = Math Processor | 6 = Video Processor |

M.14 Enumeration Values for System-System Cache-Level

| | | |
|-------------|-------------------|---------------------------|
| 1 = Other | 3 = Write Back | 5 = Varies with Address |
| 2 = Unknown | 4 = Write Through | 6 = Determination Per I/O |

M.15 Enumeration Values for System-System Cache-Cache Type

| | | |
|-------------|-----------------|-------------|
| 1 = Other | 3 = Instruction | 5 = Unified |
| 2 = Unknown | 4 = Data | |

M.16 Enumeration Values for System-System Cache-Replacement Policy

| | | |
|-----------|-------------------------------|--------------------------------|
| 1 = Other | 4 = First In First Out (FIFO) | 7 = Most Frequently Used (MFU) |
|-----------|-------------------------------|--------------------------------|

| | | |
|-------------------------------|---------------------------------|--|
| 2 = Unknown | 5 = Last In First Out (LIFO) | 8 = Data Dependent Multiple Algorithms |
| 3 = Least Recently Used (LRU) | 6 = Least Frequently Used (LFU) | |

M.17 Enumeration Values for System-System Cache-Read Policy

| | | |
|-------------|----------------|---------------------------|
| 1 = Other | 3 = Read | 5 = Read and Read-ahead |
| 2 = Unknown | 4 = Read-ahead | 6 = Determination Per I/O |

M.18 Enumeration Values for System-System Cache-Write Policy

| | | |
|-------------|-------------------|---------------------------|
| 1 = Other | 3 = Write Back | 5 = Varies with Address |
| 2 = Unknown | 4 = Write Through | 6 = Determination Per I/O |

M.19 Enumeration Values for System-System Cache-Associativity

| | | |
|-------------------|---------------------------|----------------------------|
| 1 = Other | 4 = 2-way Set-Associative | 7 = 8-way Set-Associative |
| 2 = Unknown | 5 = 4-way Set-Associative | 8 = 16-way Set-Associative |
| 3 = Direct Mapped | 6 = Fully Associative | |

M.20 Enumeration Values for System-System IRQ-Availability

| | | |
|-------------|--------------------------|------------------------------------|
| 1 = Other | 3 = Available | 5 = In Use and Available/Shareable |
| 2 = Unknown | 4 = In Use/Not Available | |

M.21 Enumeration Values for System-System IRQ-IRQ Trigger Type

| | |
|-------------|-----------|
| 1 = Other | 3 = Level |
| 2 = Unknown | 4 = Edge |

M.22 Enumeration Values for System-System DMA-Availability

| | | |
|-------------|--------------------------|--|
| 1 = Other | 3 = Available | 5 = In Use and Available/ Shareable |
| 2 = Unknown | 4 = In Use/Not Available | |

M.23 Enumeration Values for Language

The following Language enum values are displayed in the following classes: Software Group, Software Group File Information, Software, File Information, and Exclude Information.

| | | |
|--------------|---------------------------------------|--|
| 0=Neutral | 97=Not supported | 1095=Windows XP: Gujarati. This is Unicode only. |
| 1=Arabic | 101=Divehi | 1037=Hebrew |
| 2=Bulgarian | 127=Invariant Locale | 1081=Windows 2000/XP: Hindi. This is Unicode only. |
| 3=Catalan | 1024=Process or User Default Language | 1038=Hungarian |
| 4=Chinese | 2048=System Default Language | 1039=Icelandic |
| 5=Czech | 1078=Afrikaans | 1057=Indonesian |
| 6=Danish | 1052=Albanian | 1040=Italian (Standard) |
| 7=German | 1025=Arabic (Saudi Arabia) | 2064=Italian (Switzerland) |
| 8=Greek | 2049=Arabic (Iraq) | 1041=Japanese |
| 9=English | 3073=Arabic (Egypt) | 1099=Windows XP: Kannada. This is Unicode only. |
| 10=Spanish | 4097=Arabic (Libya) | 1111=Windows 2000/XP: Konkani. This is Unicode only. |
| 11=Finnish | 5121=Arabic (Algeria) | 1042=Korean |
| 12=French | 6145=Arabic (Morocco) | 2066=Windows 95 |
| 13=Hebrew | 7169=Arabic (Tunisia) | 1088=Windows XP: Kyrgyz. |
| 14=Hungarian | 8193=Arabic (Oman) | 1062=Latvian |
| 15=Icelandic | 9217=Arabic (Yemen) | 1063=Lithuanian |
| 16=Italian | 10241=Arabic (Syria) | 2087=Windows 98 only: Lithuanian (Classic) |
| 17=Japanese | 11265=Arabic (Jordan) | 1071=FYRO Macedonian |
| 18=Korean | 12289=Arabic (Lebanon) | 1086=Malay (Malaysian) |
| 19=Dutch | 13313=Arabic (Kuwait) | 2110=Malay (Brunei Darussalam) |

| | | |
|--------------------|--|---|
| 20=Norwegian | 14337=Arabic (U.A.E.) | 1102=Windows 2000/XP: Marathi. This is Unicode only. |
| 21=Polish | 15361=Arabic (Bahrain) | 1104=Windows XP: Mongolian |
| 22=Portuguese | 16385=Arabic (Qatar) | 1044=Norwegian (Bokmal) |
| 24=Romanian | 1067=Windows 2000/XP: Armenian. This is Unicode only. | 2068=Norwegian (Nynorsk) |
| 25=Russian | 1068=Azeri (Latin) | 1045=Polish |
| 26=Croatian | 2092=Azeri (Cyrillic) | 1046=Portuguese (Brazil) |
| 27=Slovak | 1069=Basque | 2070=Portuguese (Portugal) |
| 28=Albanian | 1059=Belarusian | 1094=Windows XP: Punjabi. This is Unicode only. |
| 29=Swedish | 1026=Bulgarian | 1048=Romanian |
| 30=Thai | 1109=Burmese | 1049=Russian |
| 31=Turkish | 1027=Catalan | 1103=Windows 2000/XP: Sanskrit. This is Unicode only. |
| 32=Urdu | 1028=Chinese (Taiwan) | 3098=Serbian (Cyrillic) |
| 33=Indonesian | 2052=Chinese (PRC) | 2074=Serbian (Latin) |
| 34=Ukrainian | 3076=Chinese (Hong Kong SAR, PRC) | 1051=Slovak |
| 35=Belarusian | 4100=Chinese (Singapore) | 1060=Slovenian |
| 36=Slovenian | 5124=Windows 98/Me, Windows 2000/XP: Chinese (Macau SAR) | 1034=Spanish (Spain, Traditional Sort) |
| 37=Estonian | 1050=Croatian | 2058=Spanish (Mexican) |
| 38=Latvian | 1029=Czech | 3082=Spanish (Spain, Modern Sort) |
| 39=Lithuanian | 1030=Danish | 4106=Spanish (Guatemala) |
| 41=Farsi | 1125=Windows XP: Divehi. This is Unicode only. | 5130=Spanish (Costa Rica) |
| 42=Vietnamese | 1043=Dutch (Netherlands) | 6154=Spanish (Panama) |
| 43=Armenian | 2067=Dutch (Belgium) | 7178=Spanish (Dominican Republic) |
| 44=Azeri | 1033=English (United States) | 8202=Spanish (Venezuela) |
| 45=Basque | 2057=English (United Kingdom) | 9226=Spanish (Colombia) |
| 47=FYRO Macedonian | 3081=English (Australian) | 10250=Spanish (Peru) |
| 54=Afrikaans | 4105=English (Canadian) | 11274=Spanish (Argentina) |
| 55=Georgian | 5129=English (New Zealand) | 12298=Spanish (Ecuador) |
| 56=Faeroese | 6153=English (Ireland) | 13322=Spanish (Chile) |

| | | |
|------------------|---|---|
| 57=Hindi | 7177=English (South Africa) | 14346=Spanish (Uruguay) |
| 62=Malay | 8201=English (Jamaica) | 15370=Spanish (Paraguay) |
| 63=Kazak | 9225=English (Caribbean) | 16394=Spanish (Bolivia) |
| 64=Kyrgyz | 10249=English (Belize) | 17418=Spanish (El Salvador) |
| 65=Swahili | 11273=English (Trinidad) | 18442=Spanish (Honduras) |
| 67=Uzbek | 12297=Windows 98/Me, Windows 2000/XP: English (Zimbabwe) | 19466=Spanish (Nicaragua) |
| 68=Tatar | 13321=Windows 98/Me, Windows 2000/XP: English (Philippines) | 20490=Spanish (Puerto Rico) |
| 69=Not supported | 1061=Estonian | 1072=Sutu |
| 70=Punjabi | 1080=Faeroese | 1089=Swahili (Kenya) |
| 71=Gujarati | 1065=Farsi | 1053=Swedish |
| 72=Not supported | 1035=Finnish | 2077=Swedish (Finland) |
| 73=Tamil | 1036=French (Standard) | 1114=Windows XP: Syriac. This is Unicode only. |
| 74=Telugu | 2060=French (Belgian) | 1097=Windows 2000/XP: Tamil. This is Unicode only. |
| 75=Kannada | 3084=French (Canadian) | 1092=Tatar (Tatarstan) |
| 76=Not supported | 4108=French (Switzerland) | 1098=Windows XP: Telugu. This is Unicode only. |
| 77=Not supported | 5132=French (Luxembourg) | 1054=Thai |
| 78=Marathi | 6156=Windows 98/Me, Windows 2000/XP: French (Monaco) | 1055=Turkish |
| 79=Sanskrit | 1110=Windows XP: Galician | 1058=Ukrainian |
| 80=Mongolian | 1079=Windows 2000/XP: Georgian. This is Unicode only. | 1056=Windows 98/Me, Windows 2000/XP: Urdu (Pakistan) |
| 86=Galician | 1031=German (Standard) | 2080=Urdu (India) |
| 87=Konkani | 2055=German (Switzerland) | 1091=Uzbek (Latin) |
| 88=Not supported | 3079=German (Austria) | 2115=Uzbek (Cyrillic) |
| 89=Not supported | 4103=German (Luxembourg) | 1066=Windows 98/Me, Windows NT 4.0 and later: Vietnamese |
| 90=Syriac | 5127=German (Liechtenstein) | |
| 96=Not supported | 1032=Greek | |

Setting up Security for Server Inventory

N

Server Inventory should be secured to ensure protection of all components and the database.

- ♦ The software is designed to work in a secured network, behind a firewall. Make sure all the components are within a secured network or firewall. Clients outside the firewall should connect through VPN or not connect at all.
- ♦ The database contains valuable information, that is vulnerable to hacking. Make sure the database is protected, and do not store any other data.
- ♦ Information sent over the wire media is not encrypted.
- ♦ Because log files contain information about the passwords, enable debug option only when necessary to assist in debugging.

Documentation Updates



This section contains information on documentation content changes that have been made in the *Administration* guide for Server Inventory since the initial release of Novell® ZENworks® 7 Server Management. The information will help you to keep current on updates to the documentation.

All changes that are noted in this section were also made in the documentation. The documentation is provided on the Web in two formats: HTML and PDF. The HTML and PDF documentation are both kept up-to-date with the documentation changes listed in this section.

The documentation update information is grouped according to the date the changes were published. Within a dated section, the changes are alphabetically listed by the names of the main table of contents sections for Server Inventory.

If you need to know whether a copy of the PDF documentation you are using is the most recent, the PDF document contains the date it was published on the front title page or in the Legal Notices section immediately following the title page.

The documentation was updated on the following date:

- ◆ [Section O.1, “September 07, 2007,” on page 779](#)
- ◆ [Section O.2, “July 27, 2007,” on page 780](#)
- ◆ [Section O.3, “October 19, 2006,” on page 780](#)
- ◆ [Section O.4, “August 26, 2006,” on page 780](#)
- ◆ [Section O.5, “July 14, 2006 \(Support Pack 1\),” on page 780](#)
- ◆ [Section O.6, “December 23, 2005,” on page 781](#)
- ◆ [Section O.7, “December 9, 2005,” on page 781](#)
- ◆ [Section O.8, “October 24, 2005,” on page 782](#)
- ◆ [Section O.9, “October 7, 2005,” on page 782](#)

O.1 September 07, 2007

Updates were made to the following sections. The changes are explained below.

O.1.1 Setting Up Workstation Inventory

The following updates were made in this section:

| Location | Change |
|---|---|
| Section 13.2.3, “Setting Up the MS SQL Server 2000 or MS SQL Server 2005 Inventory Database,” on page 510 | Added a warning: Do not rename the mgmtdb database. |

O.2 July 27, 2007

Updates were made to the following sections. The changes are explained below.

O.2.1 Setting Up Workstation Inventory

The following updates were made in this section:

| Location | Change |
|---|---|
| "Creating the Oracle10g Inventory Database on a UNIX Server" on page 506 | Modified the sub steps 1 and 2 in the Step 18 on page 507 to add the non-English Enum values on the Windows and Linux system. Modified the command to shutdown immediate; in the Step 21 on page 508 |
| "Creating the Oracle10g Inventory Database on a Windows Server" on page 504 | Modified the sub steps 1 and 2 in the Step 14 on page 505 to add the non-English Enum values on the Windows and Linux system. Modified the command to shutdown immediate; in the Step 17 on page 505 |

O.3 October 19, 2006

Some references to Server Inventory on Linux were removed, because it is not supported in ZENworks 7 Server Management with Support Pack 1. However, for specific information about scanning in relation to ZENworks Server Management inventoried servers, see [Section 12.2.2, "Inventory Components on Inventory Servers," on page 447](#).

O.4 August 26, 2006

Updates were made to the following sections:

| Location | Change |
|---|---|
| Appendix N, "Setting up Security for Server Inventory," on page 777 | This section is added to address security issues. |

O.5 July 14, 2006 (Support Pack 1)

Updates were made to the following sections:

| Location | Change |
|---|--|
| "Creating the Oracle10g Inventory Database on a Windows Server" on page 504 | Updated the section with Oracle10g R2 information. |

| Location | Change |
|--|--|
| "Creating the Oracle10g Inventory Database on a UNIX Server" on page 506 | Updated the section with Oracle10g R2 information. |
| Section J.3, "Performance Tips for the Inventory Server (Support Pack 1)," on page 724 | This section has been newly added. |
| Section 13.2, "Setting Up the Inventory Database," on page 493 | Updated the section with MS SQL Server 2005 information. |

O.6 December 23, 2005

Updates were made to the following sections:

| Location | Change |
|--|--|
| "Configuring the MS SQL Server 2000 Inventory Database" on page 510 | Added the following information to Step 11e on page 512 : "During the execution of the drop trigger sqls, the following error message might be displayed on the console, "Cannot drop the trigger ' <i>trigger_name</i> ', because it does not exist or you do not have permission". Ignore the error message." |
| "Configuring the Inventory Service Object" on page 520 | Newly added Step 3 on page 520 . |
| "Generating Inventory Reports" on page 660 | Added the following information as a note in Step 4 on page 661 : "ZENworks Inventory report supports only the following double-byte character languages: German, English, Spanish, French, Portuguese, and Japanese. Other double-byte characters might not be displayed properly in the Inventory reports." |
| "Sybase in the NetWare and Windows Environments" on page 717 | Added the reference to Linux in the entire section. |
| "Oracle in the NetWare, Windows, and Linux Environments" on page 719 | Added the reference to Linux in the entire section. |

O.7 December 9, 2005

Page design is reformatted to comply with revised Novell documentation standards.

O.8 October 24, 2005

Updates were made to the following sections:

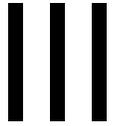
| Location | Change |
|--|--|
| "Generating Inventory Reports" on page 660 > Step 3 on page 660 | Updated the guidelines to be followed as you work with the Reporting dialog. |

O.9 October 7, 2005

Updates were made to the following sections.

| Location | Change |
|--|---|
| "Backing Up the Sybase Inventory Database" on page 498 | This section has been reorganized. There is no change in the content of the section. |
| Section 14.3.3, "Scanning for the Windows Inventoried Servers," on page 535 > "Scanning for the Hardware Inventory Information" on page 535 | Following products have been added to the list of antivirus products scanned by the Inventory scanner: Symantec AntiVirus Corporate Edition 9.0 Symantec AntiVirus Corporate Edition 10.0 |

Remote Management



The Remote Management component of Novell® ZENworks® 7 Server Management gives you the ability to manage remote servers from the management console. You can use ZENworks 7 Server Management to remotely manage Novell NetWare® 5.1/6/6.5 or Windows* 2000/2003 servers.

Remote Management can save you and your organization time and money. For example, you or your organization's help desk can analyze and remote fix server problems without having to visit the server, which reduces problem resolution times and increases productivity.

This documentation contains following sections:

- ♦ [Chapter 19, “Remote Management for NetWare Servers,” on page 785](#)
- ♦ [Chapter 20, “Remote Management for Windows Servers,” on page 795](#)
- ♦ [Appendix P, “Documentation Updates,” on page 819](#)

Remote Management for NetWare Servers

19

The Java*-based remote console utility (RConsoleJ) for Novell® ZENWorks® 7 Server Management lets you control a Novell NetWare® server and perform the following tasks:

- ♦ Use console commands as you would at the server console
- ♦ Use NLM™ programs as you would at the server console (for example, `edit.nlm` to edit files)
- ♦ Send console commands in the server's native language from the RConsoleJ Client using Buffer Input
- ♦ Control the server from another server that is using RConsoleJ
- ♦ Upgrade a NetWare server (text-based UI only)
- ♦ Secure Socket Layer (SSL) based secure session

This section contains the following topics:

- ♦ [Section 19.1, “Overview of RConsoleJ Components,” on page 785](#)
- ♦ [Section 19.2, “Setting Up RConsoleJ,” on page 786](#)
- ♦ [Section 19.3, “RConsoleJ,” on page 787](#)
- ♦ [Section 19.4, “Loading Agents at Startup,” on page 792](#)
- ♦ [Section 19.5, “Setting Up Security for RConsoleJ,” on page 792](#)
- ♦ [Section 19.6, “Managing Remote NetWare Servers,” on page 793](#)

19.1 Overview of RConsoleJ Components

RConsoleJ has the following components. These components interact with each other during the remote control session of a NetWare server.

- ♦ [Section 19.1.1, “RConsoleJ Client,” on page 785](#)
- ♦ [Section 19.1.2, “RConsoleJ Agent,” on page 785](#)
- ♦ [Section 19.1.3, “RConsoleJ Proxy Agent,” on page 786](#)

19.1.1 RConsoleJ Client

The RConsoleJ Client is a Java-based utility running on the workstation. From the RConsoleJ Client, you can remote control and monitor all NetWare console operations.

19.1.2 RConsoleJ Agent

The RConsoleJ Agent (`rconag6.nlm`) is a utility running on the target NetWare server. The target NetWare server can be connected over IP, IPX™, or IP/IPX running the RConsoleJ Agent. The RConsoleJ Agent services all RConsoleJ Client requests.

The RConsoleJ Agent advertises its services using the Service Location Protocol (SLP) on a NetWare 5.x and later.

19.1.3 RConsoleJ Proxy Agent

The RConsoleJ Proxy Agent (`rconprxy.nlm`) is a utility running on a NetWare server (supported only on Netware 5.x and up). It routes all IP packets to IPX, and vice versa.

The RConsoleJ Proxy Agent advertises its services using the Service Location Protocol (SLP).

IMPORTANT: If the target NetWare server uses only IPX, the NetWare server (loaded with the RConsoleJ Proxy Agent) must have both IP and IPX stacks installed.

19.2 Setting Up RConsoleJ

To set up RConsoleJ, complete the following sections:

- ♦ [Section 19.2.1, “Loading the RConsoleJ Agent,” on page 786](#)
- ♦ [Section 19.2.2, “Running the RConsoleJ Client,” on page 786](#)
- ♦ [Section 19.2.3, “Loading the RConsoleJ Proxy Agent on a Proxy Server,” on page 787](#)

19.2.1 Loading the RConsoleJ Agent

1 At the server console prompt, enter:

```
rconag6
```

2 Enter the password you want network administrators to use when accessing the target NetWare server using RConsoleJ.

3 Enter the TCP port number.

The default value is 2034.

If the server communicates using IPX only, enter -1 to disable TCP listening.

To enable listening over a dynamically assigned port, enter 0.

4 Enter the SPX™ port number on which RCONAG6 will listen for a proxy server.

The default is 16800.

If the server communicates using IP only, enter -1 to disable SPX listening.

To enable listening over a dynamically assigned port, enter 0.

NOTE: /DEV/TCP and /DEV/TCPSSL fail if you are using a pure IPX server.

To enable RConsoleJ across the firewall, you need to keep the following ports open: 2034, 2035, and 2036.

19.2.2 Running the RConsoleJ Client

You can run the RConsoleJ client from a workstation or a NetWare server.

- ♦ [“Running the RConsoleJ Client from a Workstation” on page 787](#)
- ♦ [“Running the RConsoleJ Client on a NetWare Server” on page 787](#)

Running the RConsoleJ Client from a Workstation

You can run the RConsoleJ Client on a workstation using any of the following methods:

- ♦ From a Windows* 2000/XP workstation, browse to *ConsoleOne_installation_directory\1.2*, and run *rconj.exe*.
- ♦ In ConsoleOne, right-click the NetWare server object that you want to remotely control, then click *Remote Management*.
- ♦ In ConsoleOne, select the NetWare server object and then from the Tools menu, click *ZENworks Remote Management > Remote Console > NetWare*.
- ♦ In ConsoleOne, right-click a Subscriber or Distributor object, then click *Remote Management*.
- ♦ If you have installed ZENworks 7 Management and Monitoring Services, select the NetWare server in the *Atlas Namespace*, then click *Remote Management*.

Running the RConsoleJ Client on a NetWare Server

You can run the RConsoleJ Client on a NetWare server using any of the following methods:

- ♦ In ConsoleOne, right-click the NetWare server object that you want to remote control and click *Remote Management*, or click the NetWare server object and then from the Tools option click *ZENworks Remote Management > Remote Console > NetWare*.
- ♦ From the server console prompt, enter *rconj.ncf*.
- ♦ From the server GUI, click *Novell > Programs > RConsoleJ*.

19.2.3 Loading the RConsoleJ Proxy Agent on a Proxy Server

The NetWare server loaded with the RConsoleJ Proxy Agent should have an IP/IPX stack loaded.

- 1 At the server console prompt, enter the following command:
`rconprxy`
- 2 Enter the TCP port number on which RCONPRXY listens for RConsoleJ.
The default is 2035.
To enable listening over a dynamically assigned port, enter 0.

When the NetWare server is running the RConsoleJ Proxy Agent, the RConsoleJ Client can communicate through it with the target NetWare server that uses only IPX to communicate.

19.3 RConsoleJ

This section will help you initiate RConsoleJ in the following scenarios:

- ♦ [Section 19.3.1, “Scenario 1: An IP Client Controlling an IP NetWare Server,” on page 787](#)
- ♦ [Section 19.3.2, “Scenario 2: An IP Client Controlling an IPX NetWare Server,” on page 790](#)

19.3.1 Scenario 1: An IP Client Controlling an IP NetWare Server

- ♦ [“Prerequisites” on page 788](#)
- ♦ [“Setting Up a Secured IP Connection” on page 788](#)

- ◆ “Setting Up an Unsecure IP Connection” on page 789

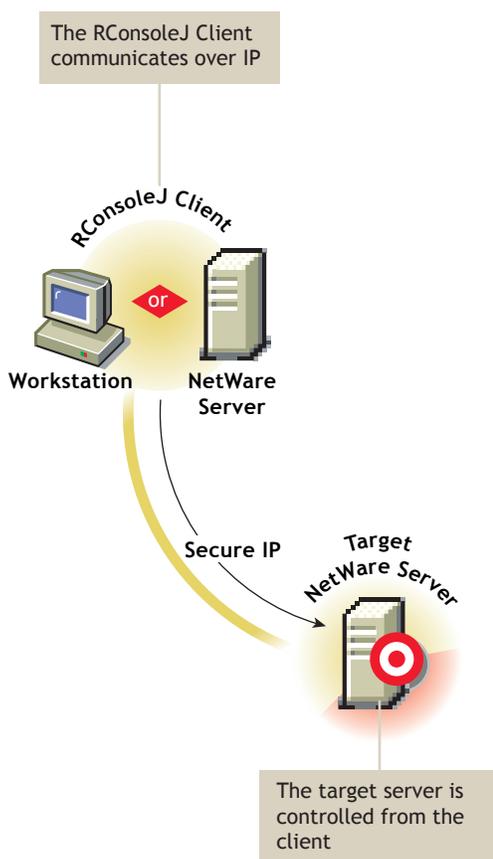
Prerequisites

- ◆ “Loading the RConsoleJ Agent” on page 786
- ◆ “Running the RConsoleJ Client” on page 786

Setting Up a Secured IP Connection

Figure 19-1 illustrates how the RConsoleJ Client communicates directly with the RConsoleJ Agent using TCP/IP.

Figure 19-1 RConsoleJ Client Communicates with the Target NetWare Server Over Secure IP



To setup a Secured IP connection, complete the following:

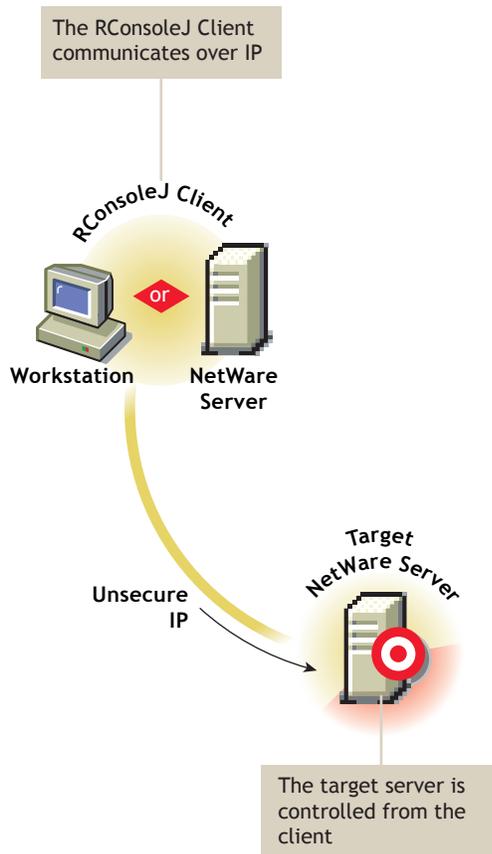
- 1 In the RConsoleJ Connection dialog box, select *Secure IP*.
- 2 Enter the IP address of the target NetWare server, or click the *Remote Servers* icon and then select the target NetWare server from the list.
- 3 Enter the password provided during loading of the RConsoleJ Agent on the target server.
- 4 Enter the port number.
The default is 2036.
- 5 Click *Connect*.

- 6 To ensure server authentication, read the Untrusted Certificate Verification server certificate issued by the target server and click *OK* to accept.

Setting Up an Unsecure IP Connection

Figure 19-2 illustrates how the RConsoleJ Client communicates directly with the RConsoleJ Agent using TCP/IP.

Figure 19-2 RConsoleJ Communicates with the Target NetWare Server Over IP



When you run the RConsoleJ client from ConsoleOne, the Novell RConsoleJ dialog box is displayed with the Netware Server IP address. To run the RConsoleJ client, see [“Running the RConsoleJ Client” on page 786](#).

To start an IP connection:

- 1 Enter the password specified during loading the RConsoleJ Agent.
- 2 Enter the port number.
The default is 2034.
- 3 Click *Connect*.

19.3.2 Scenario 2: An IP Client Controlling an IPX NetWare Server

- ♦ “Prerequisites” on page 790
- ♦ “Starting an IPX Connection” on page 790

Prerequisites

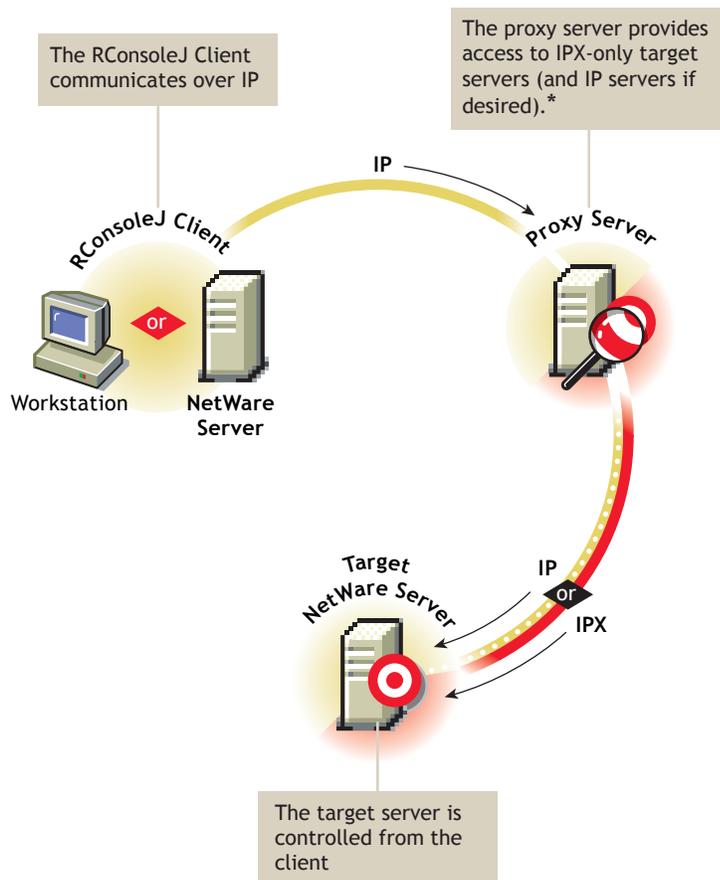
- ♦ “Loading the RConsoleJ Agent” on page 786
- ♦ “Running the RConsoleJ Client” on page 786
- ♦ “Loading the RConsoleJ Proxy Agent on a Proxy Server” on page 787

Starting an IPX Connection

The RConsoleJ Client communicates with the RConsoleJ Agent through the RConsoleJ Proxy Agent because the target NetWare server is based only on IPX.

The RConsoleJ Proxy Agent is loaded on a NetWare server (proxy server) that has both IP and IPX stacks loaded. The RConsoleJ Proxy Agent receives all the IP requests from the RConsoleJ Client, converts them to IPX requests, and then sends them to the RConsoleJ Agent and vice-versa. [Figure 19-3](#) illustrates this.

Figure 19-3 The RConsoleJ Client Communicates with the Target NetWare Server through the Proxy Server



*If the target server uses IPX, the proxy server must have both IP and IPX stacks loaded.

When you run the RConsoleJ client from ConsoleOne, the Novell RConsoleJ dialog box is displayed with the Netware Server IP address. To run the RConsoleJ client, see [“Running the RConsoleJ Client” on page 786](#).

To start an IPX connection:

- 1 From the Connect Type drop-down list, select *Connect through Proxy*. Select *SPX* to get the IPX address and *TCP* to get the IP address.

The default port is be selected when you make the above change.

The default is 16800 for IPX address and 2034 for IP address.

- 2 Enter the IP address of the proxy server, or click the *Remote Servers* icon and then select a proxy server from the list.

- 3 Enter the port number specified during loading the RConsoleJ Proxy Agent.

The RConsoleJ Client communicates with the RConsoleJ Proxy Agent on this port.

The default is 2035.

- 4 Click *Connect*.

19.4 Loading Agents at Startup

You can load the RConsoleJ Agent and the RConsoleJ Proxy Agent at the startup.

- [Section 19.4.1, “Loading RConsoleJ Agent at Startup,” on page 792](#)
- [Section 19.4.2, “Loading RConsoleJ Proxy Agent at Startup,” on page 792](#)

19.4.1 Loading RConsoleJ Agent at Startup

You can load RConsoleJ Agent (`rconag6.nlm`) at startup using either of the following methods:

- [“Using Encrypted Password \(Recommended\)” on page 792](#)
- [“Using Plain Text Password” on page 792](#)

Using Encrypted Password (Recommended)

- 1 Prepare the `ldrconag.ncf` script file.
 - 1a At the console prompt, enter `rconag6 encrypt`.
 - 1b In the RConsoleJ Agent server screen, enter the password and port numbers.
 - 1c Enter `Y`.
- 2 Prepare the `autoexec.ncf` file to run the `ldrconag.ncf` script file at startup.
 - 2a Comment out the following line:

```
load rconag6 user_defined_password_here 2034 16800 2036
```
 - 2b Enter the following line:

```
ldrconag
```

NOTE: The ZENworks 7 Server Management Remote Management Agent install automatically executes the above steps if you specified a password for the managed server in the installation wizard.

Using Plain Text Password

Enter the following line in `autoexec.ncf`:

```
load rconag6 your_password 2034 16800 2036
```

19.4.2 Loading RConsoleJ Proxy Agent at Startup

To load the RConsoleJ Proxy Agent (`rconprxy.nlm`) at the startup, enter the following line in `autoexec.ncf`:

```
load rconprxy 2035
```

19.5 Setting Up Security for RConsoleJ

You can change the agent password to ensure that RConsoleJ sessions are secure.

To change the agent password for a remotely managed NetWare server:

- 1 At the NetWare Console prompt, enter `unload rconag6` to unload `rconag6.nlm`.

- 2 Enter `rconag6 encrypt`.
- 3 Enter a new password.
- 4 Enter the TCP port number.
The default is 2034.
- 5 Enter the SPX port number.
The default is 16800.
- 6 Enter `y` when prompted to save the following command line in the `ldrconag.ncf` file.
If you enter `n`, the `ldrconag.ncf` file is not updated. The new password is valid only for the current session. If you load RCONAG6 from the `ldrconag.ncf` file later, the previously saved password is used.

The new password is in effect when the agent is loaded from the LDRCONAG script file after you reboot the server.

Alternatively, you can change the password using the following commands:

- ♦ For Netware 4.x and 5.x:

```
LOAD RCONAG6 - E encrypted_password TCP_port_number  
SPX_port_number
```

- ♦ For Netware 6.x:

```
LOAD RCONAG6 - E encrypted_password TCP_port_number  
SPX_port_number Secure_port_number
```

19.6 Managing Remote NetWare Servers

After RConsoleJ establishes connection with the NetWare server, you can view and manage the target NetWare server from your desktop.

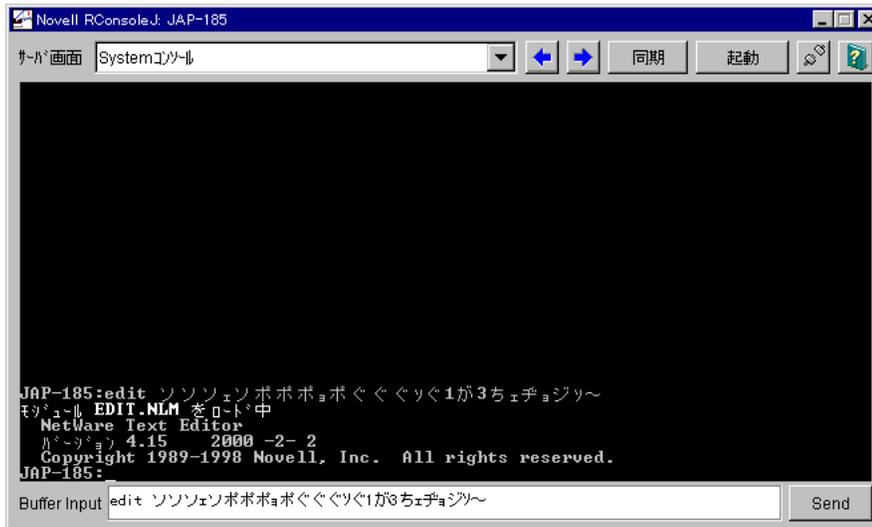
The following sections explain the tasks you can perform to effectively manage a remote NetWare Server:

- ♦ [Section 19.6.1, “Sending Console Commands in the Server's Native Language,”](#) on page 793
- ♦ [Section 19.6.2, “Synchronizing RConsoleJ Client and Target NetWare Screens,”](#) on page 794

19.6.1 Sending Console Commands in the Server's Native Language

You can send console commands in the server's native language from the RConsoleJ Client using the *Buffer Input* field as shown in [Figure 19-4](#). The buffer stores a list of ten history commands.

Figure 19-4 Send Console Commands in Japanese using Buffer Input



To send console commands, do the following:

- 1 From the Novell RConsoleJ Client window, enter the command that you want to run at the target server in the *Buffer Input* field.
- 2 Click *Send*.

19.6.2 Synchronizing RConsoleJ Client and Target NetWare Screens

You can synchronize the screen displayed on the target NetWare server and the screen displayed on the RConsoleJ Client with each other. Switching the screen in the RConsoleJ Client switches the screen at the target server console, and vice versa.

To synchronize the screens, from the Novell RConsoleJ window, click *Sync*.

To switch the screen on the server console to the currently activated screen on the RConsoleJ Client, click *Activate*.

Remote Management for Windows Servers

20

The Remote Management component of Novell® ZENworks® 7 Server Management allows you to remotely manage Windows* 2000/2003 servers from your computer.

This chapter contains the following topics:

- ♦ [Section 20.1, “Remote Management Terminology,” on page 795](#)
- ♦ [Section 20.2, “Understanding Remote Management for Windows Servers,” on page 795](#)
- ♦ [Section 20.3, “Setting Up Security for Remote Management,” on page 797](#)
- ♦ [Section 20.4, “Managing Remote Windows Servers,” on page 800](#)

20.1 Remote Management Terminology

The following brief glossary provides basic definitions of Remote Management terms:

Managed server: A Windows server that you want to remotely manage. You must install the ZENworks 7 Remote Management Agent on it. If you manage the server through the Server Remote Management Policy, you must install the ZENworks 7 Subscriber component also.

Management console: A Windows 2000/XP workstation or 2000/2003 server running Novell ConsoleOne® with the ZENworks 7 Remote Management ConsoleOne snap-ins installed. The management console provides the interface where you manage and administer your network.

Management server: A server with Novell eDirectory™ and the ZENworks 7 Distributor components. The eDirectory and Distributor components must be installed if you want to manage servers through the Server Remote Management Policy. Your management server can be a managed server.

Remote operator: A user who can remotely manage servers.

Administrator: A person who has the rights to install Remote Management components. All administrators are remote operators but all remote operators are not administrators.

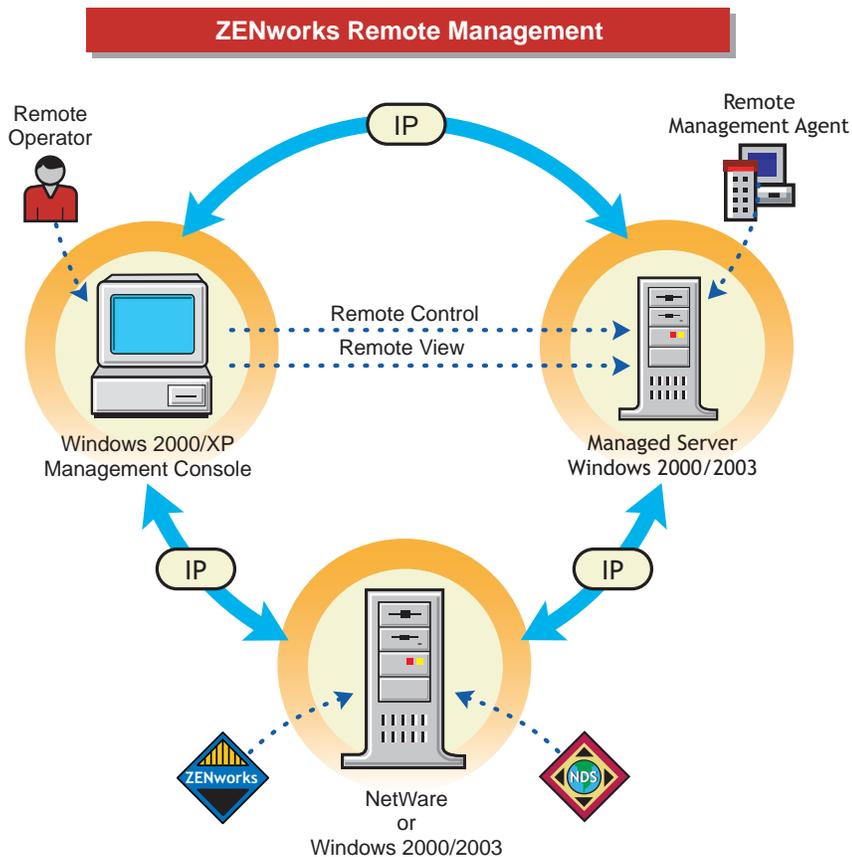
Remote Management Agent: A Server Management component that is installed on a managed server so the remote operator can remotely manage that server. The Remote Management Agent starts automatically when the managed server boots up and authenticates the remote operator when the remote session is initiated.

Viewing window: A representation of the managed server desktop. It is displayed on the management console when the remote operator initiates a Remote Management session.

20.2 Understanding Remote Management for Windows Servers

Figure 20-1 depicts the functionality of the ZENworks 7 Remote Management, which is explained below:

Figure 20-1 Remote Management Functionality



To remotely manage a server, the ZENworks 7 Remote Management Agent component must be installed on that server. For more information, see “[Policy-Enabled Server Management Installation](#)” in the *Novell ZENworks 7 Server Management Installation Guide*.

The Remote Management Agent starts automatically when the managed server boots up. The agent password can be set either by the administrator during the Remote Management installation or by the user at the managed server after the installation. The remote operator would be required to enter the password when he or she initiates a Remote Management session with a managed server. On successful verification, the Remote Management session proceeds and the Viewing window is displayed on the management console.

The user at the managed server can change the agent password. In such a case, the new password will have to be communicated to the remote operator. It is recommended that the password be changed frequently to prevent it from being compromised. For more details on how to set up a new agent password or change the existing agent password, see “[Setting Up the Agent Password at the Managed Server](#)” on page 800.

20.3 Setting Up Security for Remote Management

The information in the following sections help you in setting up security for the Remote Management sessions:

- ♦ [Section 20.3.1, “Configuring the Remote Management Policies,” on page 797](#)
- ♦ [Section 20.3.2, “Setting Up the Agent Password at the Managed Server,” on page 800](#)

20.3.1 Configuring the Remote Management Policies

To configure the Remote Management policies, you must perform the following tasks:

- ♦ [“Creating the Policy Packages” on page 797](#)
- ♦ [“Creating and Configuring the Tiered Electronic Distribution Objects” on page 798](#)
- ♦ [“Configuring the Server Remote Management Policy” on page 798](#)
- ♦ [“Configuring the Distribution Object for Remote Management” on page 800](#)
- ♦ [“Configuring the Distributor and the Subscriber Objects” on page 800](#)

You can also change the security settings on the managed servers by modifying the [Remote Management Policy] section in the

`ZENworks_agent_directory\rmagent\zfsrpol.ini` file.

Creating the Policy Packages

ZENworks 7 requires policy packages in the eDirectory tree that can hold the server policies. You can later configure and enable the server policies.

Policy packages are eDirectory objects that contain collections of policies grouped according to the object types. You should create an Organizational Unit (OU) for holding the policy packages.

Consider the following when determining where to place this OU:

- ♦ Whether you have partitions in your tree
- ♦ The 256-character limit in eDirectory for the full distinguished name
- ♦ How you will use the Search policy to locate the policy package

If you install ZENworks 7 Desktop Management to your tree, you may want to keep the ZENworks Server Management and Desktop Management policies in separate containers, such as `Server_Policies` and `Desktop_Policies`.

For Remote Management, create two containers, one for Tiered Electronic Distribution objects and the other for the Remote Management policy package.

To create a container:

- 1 In ConsoleOne, right-click the container where you want the container for the policy packages placed.
- 2 Click *New > Object > Organizational Unit > OK*.
- 3 Name the container, for example, `Server_Policies`, then click *OK*.

IMPORTANT: If you have partitions that are accessed across a WAN, make sure that the Policy Package objects are in the same partition as the Server object so that the Policy/Package Agents are loaded. Also make sure that the Search policy does not require searching outside the partition where the Server object exists.

For Remote Management, you must create the Distributed Server package. The Distributed Server package is required to distribute the Remote Management policies among the managed servers for enforcement.

To create the Distributed Server package:

- 1 In ConsoleOne, right-click the policy package's container, then click *New > Policy Package*.
The Policy Package Wizard is displayed.
- 2 In the *Policy Packages* list, select *Distributed Server Package*, then click *Next*.
- 3 Enter a name for the Distributed Server Package, then click *Next*, then click *Finish*.

Creating and Configuring the Tiered Electronic Distribution Objects

For Remote Management, you must create and configure the following Tiered Electronic Distribution objects:

- ◆ TED Distribution
- ◆ TED Channel

To create and configure the Tiered Electronic Distribution objects, see [Chapter 3, “Tiered Electronic Distribution,” on page 85](#).

Configuring the Server Remote Management Policy

The Server Remote Management Policy defines the behavior of the Remote Management Agent. This policy is distributed to the specified Windows managed servers using the Tiered Electronic Distribution, which helps the remote operator to associate the Remote Management policy to a group of Windows managed servers from the management console.

To configure the Server Remote Management Policy:

- 1 In ConsoleOne, right-click the Distribute Server Package object, then click *Properties*.
- 2 Click the *Policies* tab and select the *Windows* sub-option.
- 3 Select the check box under the *Enabled* column for the Server Remote Management Policy.
- 4 Click the *Properties* button > the *Remote Management* tab.
- 5 Click the *General* tab, then select the any of following options:
 - ◆ **Enable Session Encryption:** Encrypts the Remote Control and Remote View sessions. The Remote Operator cannot change this to an unencrypted mode. If you do not select this check box, the remote sessions are unencrypted by default. In this case, the Remote Operator has an option to switch over to the encrypted mode from the Console. An encrypted session slightly impacts the performance of remote sessions over fast links.

IMPORTANT: This option does not work for ZENworks for Servers 3.x and earlier versions of the Remote Management Agent.

- ♦ **Allow User to Request Remote Session:** Enables the user at the managed server to request the Remote Operator on the management console to perform a remote session.

IMPORTANT: This option does not work for ZENworks for Servers 3.x and earlier versions of the Remote Management Agent.

- ♦ **Display Remote Management Agent Icon To Users:** Displays the *Remote Management Agent* icon in the system tray of the Windows 2000 or Windows 2003 managed servers on which the Remote Management Agent is running.

6 Click the *Remote Control* tab, then select the any of following options:

- ♦ **Prompt User for Permission to Remote Control:** Allows the user at the managed server to either accept or reject the Remote Control session initiated by the remote operator.
- ♦ **Give User Audible Signal when Remote Controlled:** Generates an audible signal on the managed server every time the remote operator remote controls the managed server. You can modify the time interval as to when you want the audible signal should be generated.
- ♦ **Give User Visible Signal when Remote Controlled:** Displays a visible signal with the name of the remote operator and console machine on the managed server every time the remote operator remote controls the managed server. You can modify the time interval as to when the name should be displayed.
- ♦ **Allow Blanking User's Screen:** Allows the remote operator to blank the screen of the managed server during a remote control session and also locks the mouse and keyboard controls.
- ♦ **Allow Locking User's Keyboard and Mouse:** Allows the remote operator to lock the keyboard and mouse controls of the managed server during a remote control session.

7 Click the *Remote View* tab, then select the any of following options:

- ♦ **Prompt User for Permission to Remote View:** Allows the user at the managed server to either accept or reject the Remote View session initiated by the remote operator.
- ♦ **Give User Audible Signal when Remote Viewed:** Generates an audible signal on the managed server every time the remote operator remotely views the managed server. You can modify the time interval as to when you want the audible signal should be generated.
- ♦ **Give User Visible Signal when Remote Viewed:** Displays a visible signal with the name of the remote operator and console machine on the managed server every time the remote operator remotely views the managed server. You can modify the time interval as to when the name should be displayed.

8 Click *Apply*, then click *Close*.

9 Right-click the Server Remote Management Policy, then select *Edit Schedule*.

10 Modify the schedule.

11 Click *Apply*, then click *Close*.

12 To associate the Server Remote Management Policy with a managed server, click the *Distribution* tab.

13 Click *Add*.

14 Browse for and select the Distribution object, then click *OK*.

15 Click *Apply*, then click *Close*.

Configuring the Distribution Object for Remote Management

You must configure the Distribution object for distributing the Remote Management policies.

To configure the Distribution object:

- 1 In ConsoleOne, right-click the Distribution object, then click *Properties*.
- 2 Click the *Type* tab.
- 3 Select Policy Package from the *Select Type* drop-down list.
- 4 Click *Add*, then select the Distributed Server package that has the Server Remote Management Policy.
- 5 Click the *Schedule* tab.
- 6 Modify the schedule.
- 7 Click *Apply*, then click *Close*.

Configuring the Distributor and the Subscriber Objects

To configure the Distributor and the Subscriber objects, see [Chapter 3, “Tiered Electronic Distribution,” on page 85](#).

If the managed servers are residing on a different eDirectory tree or the Windows 2000/2003 server does not have the eDirectory installed, you must create and configure an External Subscriber object for sending Distributions to Subscribers residing on managed servers in other trees. For more information on External Subscribers, see [Section 3.8, “External Subscribers,” on page 156](#).

20.3.2 Setting Up the Agent Password at the Managed Server

The user at the managed server can change the password of the Remote Management Agent to make sure that the Remote Management sessions are secure.

To change the agent password:

- 1 Right-click the *Remote Management Agent* icon from the system tray of the Windows 2000/2003 managed server.
- 2 Click *Security > Set Password*.
Use a password of ten or fewer ASCII (non-extended) characters. The password is case-sensitive and cannot be blank.

The new password must be communicated to the remote operator each time it is changed.

20.4 Managing Remote Windows Servers

The following sections provide information that will help you effectively manage Remote Management sessions on Windows 2000/2003 servers:

- ◆ [Section 20.4.1, “Initiating Remote Management Sessions,” on page 801](#)
- ◆ [Section 20.4.2, “Operating with Windows XP SP2,” on page 803](#)
- ◆ [Section 20.4.3, “Configuring Remote Management Ports,” on page 803](#)
- ◆ [Section 20.4.4, “Customizing the Permission Message,” on page 804](#)

- ◆ Section 20.4.5, “Managing a Remote View Session,” on page 805
- ◆ Section 20.4.6, “Managing a Remote Control Session,” on page 807
- ◆ Section 20.4.7, “Remote Operator Identification Display,” on page 815
- ◆ Section 20.4.8, “Viewing the Audit Log for Remote Management Sessions,” on page 815
- ◆ Section 20.4.9, “Improving the Remote Management Performance,” on page 815
- ◆ Section 20.4.10, “Shutting Down and Restarting the Remote Management Agent,” on page 816

20.4.1 Initiating Remote Management Sessions

You have several options for initiating a Remote Management session from ConsoleOne. They include the following:

- ◆ “Initiating Remote Management Session from the ConsoleOne Tools Menu” on page 801
- ◆ “Initiating Remote Management Session from the eDirectory/NDS Namespace” on page 801
- ◆ “Initiating Remote Management Session from the Atlas Namespace” on page 802
- ◆ “Initiating Remote Management Session from the Remote Management Agent” on page 802

Initiating Remote Management Session from the ConsoleOne Tools Menu

- 1 In ConsoleOne, click *Tools > ZENworks Remote Management > Remote Console > Windows*.
- 2 In the Remote Management dialog box, enter the IP address or the DNS name of the managed server.
- 3 Enter the agent password.
- 4 Select the Remote Management operation that you want to initiate with the managed server.
- 5 Click *OK*.

Initiating Remote Management Session from the eDirectory/NDS Namespace

You can start a Remote Management session from the eDirectory (NDS) namespace (in ConsoleOne) using one of the following methods:

- 1 In ConsoleOne, select a managed server.
- 2 Click *Tools > ZENworks Remote Management > Remote Console > Windows*.
- 3 In the Remote Management dialog box, select the IP address of the managed server from the *Agent* drop-down list.
The IP address of the selected managed server is automatically populated to the *Agent* drop-down list.
- 4 Enter the agent password.
- 5 Select the Remote Management operation that you want to initiate with the managed server.
- 6 Click *OK*.

You can also use the following procedure:

- 1 In ConsoleOne, right-click a managed server.
- 2 Click *Remote Management*.

- 3 In the Remote Management dialog box, select the IP address of the managed server from the *Agent* drop-down list.
The IP address of the selected managed server is automatically populated to the *Agent* drop-down list.
- 4 Enter the agent password.
- 5 Select the Remote Management operation that you want to initiate with the managed server.
- 6 Click *OK*.

Initiating Remote Management Session from the Atlas Namespace

Before initiating a Remote Management session from the Atlas namespace (in ConsoleOne), make sure that the NetWare[®] Management Agent[™] (NMA) is installed and the Discovery discovers the network topology.

To initiate the Remote Management session:

- 1 In ConsoleOne, right-click a managed server.
- 2 Click *Actions > Remote Control* or *Remote View*.
- 3 Select the IP address and enter the agent password.
The IP address of the selected managed server is automatically populated to the *Agent* drop-down list.
- 4 Click *OK*.

Initiating Remote Management Session from the Remote Management Agent

If the managed server is configured behind dynamic NAT, the managed server cannot be accessed from the management console but the management console can be accessed from the managed server. To resolve this problem:

- 1 The user at the managed server must initiate a request for a Remote Management session to the remote operator by using the Request Session option.

IMPORTANT: Before initiating a Remote Management session from the Remote Management Agent, the remote operator must ensure that ConsoleOne is running on the management console.

NOTE: The first instance of ConsoleOne receives the request when a session request is initiated from a Remote Management Agent to the management console running on a terminal server. None of the ConsoleOne instances receive the session request until all ConsoleOne instances are closed on the session where ConsoleOne was launched for the first time. To receive the session request, ConsoleOne must be launched again on any terminal session.

To request for a session, the user at the managed server must do the following:

- 1a Right-click the *Remote Management Agent* icon.
- 1b Select *Request Session*.
- 1c Enter the IP address or the DNS name of the management console.
- 1d Select the *Remote Control* or *Remote View* operation from the drop-down list.

- 1e** Click *OK*.
- 2** The Remote Management Listener listens to the request and notifies the remote operator about it. The remote operator must accept the request and provide the following credentials for the request in the Select Authentication Mode dialog box:
 - 2a** Enter the password for authentication.
 - 2b** Click *OK*.

20.4.2 Operating with Windows XP SP2

Windows XP SP2 comes with a firewall enabled by default. As a result, the Remote Control Listener running on Windows XP SP2 cannot receive connections initiated by the Remote Management Agent.

You need to configure the firewall settings to allow the Remote Control Listener to receive connections.

The Remote Control Listener binds to TCP port 1762 by default. In order to change the ports, refer to [“Configuring Remote Control Listener Port” on page 804](#).

20.4.3 Configuring Remote Management Ports

This section provides information on the following topics:

- ◆ [“Configuring Remote Management Agent Port” on page 803](#)
- ◆ [“Configuring Remote Control Listener Port” on page 804](#)

Configuring Remote Management Agent Port

The Remote Management Agent port binds to TCP port 1761 by default. You may configure it to run on a different TCP port by following the steps mentioned below:

- 1** Open `ZENworks_agent_directory\rmagent\rmcfg.ini` file.
- 2** In the *Remote Management Agent Port* section, set the *DefaultCommPort* to the desired port number.
- 3** Restart the Novell ZENworks Remote Management Agent service.

To initiate a remote session to a managed server where the Remote Management Agent is running on any port other than 1761, do the following modifications on the management console:

- 1** Open the `ConsoleOne_directory\1.2\bin\rmports.ini` file.
- 2** In the *Remote Management Agent Ports* section, add the port number.

NOTE: If the Remote Management Agents are running on different ports on different managed servers, you may mention the port numbers one below the other under the Remote Management Agent Ports section.

Configuring Remote Control Listener Port

The Remote Control Listener port binds to TCP port 1762 by default when ConsoleOne is started. You may configure it to run on a different TCP port by following the steps mentioned below:

- 1 Open the *ConsoleOne_directory\1.2\bin\rmports.ini* file.
- 2 In the *Remote Control Listener Port* section, set the *DefaultCommPort* to the desired port number.
- 3 Restart ConsoleOne.

To initiate a remote session request to a management console, where the Remote Control Listener is running on any port other than 1762, the following modifications need to be done on the managed servers:

- 1 Open the *ZENworks_agent_directory\rmagent\rmcfg.ini* file.
- 2 In the *Remote Control Listener Ports* section, add the port number.

NOTE: If the Remote Control Listeners are running on different ports on different management consoles, you may mention the port numbers one below the other in the *Remote Control Listener Ports* section.

20.4.4 Customizing the Permission Message

If the *Ask for user permission* option is selected in the Remote Management policy, the Request for Permission dialog box is displayed with the following default message:

Do you want to allow console user to perform remote management operation?

ZENworks 7 with Support Pack 1 allows you to customize the default message displayed in the Request for Permission dialog box.

To customize the default message, do the following on the managed server:

- 1 Open the Registry Editor.
- 2 Traverse to `HKEY_LOCAL_MACHINE\Software\Novell\ZENworks\RemoteManagement\RMAgent` and create a registry string in the name "PermissionMessage".
- 3 Enter the message that should be displayed in the Request for Permission dialog box as the value of the registry string created in the previous step.
- 4 (Optional) In the registry string value, you can use the following parameters that will be dynamically replaced by valid information in the message:

Table 20-1 Parameters Used to Customize the Message of the Request for Permission dialog box

| Parameter | Information Displayed |
|-----------|--|
| %a or %A | Displays the Remote console user name. |
| %i or %I | Displays the IP address of the management console. |
| %r or %R | Displays the Remote Management operation initiated by Remote Operator. |

A sample registry string with parameters is as follows:

```
Do you want to allow %a to %r from the remote machine, %i?
```

The registry string is displayed as the following message in the Request for Permission dialog box:

```
Do you want to allow admin.novell to Remote Control from the remote machine,10.0.0.0?
```

20.4.5 Managing a Remote View Session

After you have initiated a Remote Management session and selected Remote View as the operation, you have several options to help you view the managed server.

- ◆ [“Controlling the Display of the Viewing Window” on page 805](#)
- ◆ [“Using the Viewing Window Accelerator Keys” on page 806](#)
- ◆ [“Defining a Custom Accelerator Key Sequence” on page 807](#)

Controlling the Display of the Viewing Window

You can regulate the display of the Viewing window through using the control options.

To enable the control options:

- 1 Click the *Remote Management Agent* icon, located at the top left corner of the Viewing window.
- 2 Click *Configure*.

| Option | Description |
|------------------------------------|---|
| <i>Enable High Quality Scaling</i> | Enhances the quality of images in the Scale To Fit Mode. |
| <i>Enable Accelerator Keys</i> | Enables the accelerator keys on the management console so that you can change the default accelerator key sequences during the remote session. |
| <i>Enable Encryption</i> | Encryption is an optional feature and will be effective per session. If the saved configuration has enabled encryption, the session will be encrypted from the start of the session. Encrypting a whole session provides greater security as the data transferred over the wire will be encrypted and it will be difficult to decipher anything meaningful even after the data over the wire is captured. However, it impacts performance slightly and is recommended when the security requirement is very stringent. |
| <i>Hide Wallpaper</i> | Suppresses any wallpaper displayed on the managed server. This option is enabled by default. If you want to display the wallpaper on the managed server during a Remote View session, disable this option. |

| Option | Description |
|----------------------|---|
| <i>Color Quality</i> | <p>By default, on a fast Link, the color quality is set to Normal and on a slow link the color quality is set to 256 colors. You can change the color quality of the slow link or the fast link to one of the following:</p> <ul style="list-style-type: none"> ◆ 16 Colors: Forces the use of 16-color palette on the managed server during a Remote Management session. This enhances the Remote Management performance particularly over a slow-link. ◆ 256 Colors: Forces the use of 256-color palette on the managed server during a Remote Management session. This enhances the Remote Management performance over a slow-link. ◆ Normal: The color is not altered and the setting is the same on the managed server during a Remote Management session. |
| <i>Network Type</i> | <p>if the managed server is connected by a LAN, select the <i>Fast Links</i> option to enhance the Remote Management performance.</p> <p>if he managed server is connected over a dial-up link or by WAN, select the <i>Slow Links</i> option to enhance the Remote Management performance.</p> |

- 3 To save the *Control Parameter* settings, select the *Save on Exit* check box.
The saved settings are implemented in the next Remote View session.
- 4 Click *OK*.

Using the Viewing Window Accelerator Keys

You can use accelerator keys to assign the shortcut keys to the control options and also to control the display of the Viewing window. Default accelerator key sequences are assigned to each accelerator key option. The Accelerator Keys dialog box displays the default key sequence in the edit field of each accelerator key option. You can define a custom accelerator key sequence to change the default sequence. For more information, see [“Defining a Custom Accelerator Key Sequence” on page 807](#).

To enable the Accelerator Keys option:

- 1 Click the *Remote Management Agent* icon, located at the top left corner of the Viewing window.
- 2 Click *Configure*.
- 3 Select *Enable Accelerator Keys*.
- 4 Click *OK*.

To open the Accelerator Keys dialog box:

- 1 Click the *Remote Management Agent* icon, located at the top left corner of the Viewing window.
- 2 Click *Accelerator Keys*.

Table 20-2 explains the Accelerator Key options you can during the Remote View session:

Table 20-2 Accelerator Key Options for Remote View Session

| Option | Default Keystroke | Description |
|--------------------------------|-------------------|--|
| <i>Toggle Full Screen</i> | Ctrl+Alt+M | Applicable only if the color resolution settings on the management console and managed server are same. Sizes the Viewing window to the size of your screen without window borders. |
| <i>Refresh Screen</i> | Ctrl+Alt+R | Refreshes the Viewing window. |
| <i>Restart Session</i> | Ctrl+Alt+T | Re-establishes the connection with the managed server. |
| <i>Enable Accelerator Keys</i> | Ctrl+Alt+A | Allows you to enable or disable the default accelerator key sequences. |
| <i>Stop Viewing</i> | Left-Shift+Esc | Closes the Viewing window. |
| <i>Configure Dialog</i> | Alt+M | Opens the Control Parameters dialog box. |
| <i>Accelerator Keys Dialog</i> | Alt+A | Opens the Accelerator Keys dialog box. |
| <i>Poll Full Screen</i> | Alt + L | Scans and renders the information of the entire screen of the managed server. |
| <i>Scale To Fit</i> | Ctrl+Alt+G | Hides the scroll bars and scale the Remote Management window to fit your screen. |

Defining a Custom Accelerator Key Sequence

The default keystrokes assigned to the accelerator key options are displayed in the edit field to the right of each accelerator key option in the Accelerator Keys dialog box. You can change the accelerator key sequence and define a custom accelerator key sequence if you do not want to use the default keystroke.

To define a custom accelerator key sequence:

- 1 Click the *Remote Management Agent* icon, located at the top-left corner of the Viewing window.
- 2 Click *Accelerator Keys*.
- 3 Click the *Edit* field of the accelerator key option where you want to define a custom accelerator key sequence.
- 4 Press the new accelerator key sequence.
- 5 Click *OK*.

IMPORTANT: The shift keys are left-right sensitive, and are indicated in the Control Options dialog box as LShift and RShift. Avoid the use of standard key sequences like Ctrl+C, Ctrl+V, Shift+Del, etc.

20.4.6 Managing a Remote Control Session

After you have initiated a Remote Management session and selected Remote Control as the operation, you can control the managed server from the management console to provide user

assistance and to help resolve server problems. With remote control connections, the remote operator can go beyond viewing the managed server to taking control of it.

You can effectively manage a Remote Control session by performing the following tasks with the Viewing window control options, the Viewing window toolbar buttons, and the Remote Management Agent icon options:

- ◆ “Controlling the Display of the Viewing Window” on page 808
- ◆ “Using the Viewing Window Accelerator Keys” on page 809
- ◆ “Using the Toolbar Buttons on the Viewing Window” on page 811
- ◆ “Enabling the Wallpaper on the Managed Server” on page 812
- ◆ “Using the Remote Management Agent Icon” on page 813
- ◆ “Setting Up a Password for the Managed Server” on page 813
- ◆ “Obtaining Information About Remote Management Sessions” on page 814
- ◆ “Obtaining General Information” on page 814
- ◆ “Obtaining Security Information” on page 814

Controlling the Display of the Viewing Window

You can control the display of the managed server by using the Viewing window control options.

To enable control options:

- 1 Click the *Remote Management Agent* icon, located at the top left corner of the Viewing window.
- 2 Click *Configure*.
- 3 Select the control options you want to enable for the remote session.

The following table explains the options you can use to control the display of the Viewing window.

| Option | Description |
|---------------------------------------|---|
| <i>Block Mouse Movements to Agent</i> | To reduce network bandwidth consumption, blocks all the mouse movements to the Agent. |
| <i>Enable High Quality Scaling</i> | Enhances the quality of images in the Scale To Fit mode. |
| <i>Enable Accelerator Keys</i> | Enables the accelerator keys on the management console so that you can change the default accelerator key sequences during the remote session. |
| <i>Enable Encryption</i> | Encryption is an optional feature and will be effective per session. If the saved configuration has enabled encryption, the session will be encrypted from the start of the session. Encrypting a whole session provides greater security as the data transferred over the wire will be encrypted and it will be difficult to decipher anything meaningful even after the data over the wire is captured. However, it impacts performance slightly and is recommended when the security requirement is very stringent. |

| Option | Description |
|--------------------------|---|
| <i>System Key Pass</i> | <p>Passes Alt-key sequences on the management console to the managed server.</p> <hr/> <p>NOTE: During a Remote View session, the <i>System Key pass Through</i> option is not enabled.</p> |
| <i>Hide Wallpaper</i> | <p>Suppresses any wallpaper displayed on the managed server. This option is enabled by default. If you want to display the wallpaper on the managed server during a Remote Control or Remote View session, disable this option.</p> |
| <i>Enable Encryption</i> | <p>Encryption is an optional feature and will be effective per session. If the saved configuration has enabled encryption, the session will be encrypted from the start of the session.</p> <p>Encrypting a whole session provides greater security as the data transferred over the wire will be encrypted and it will be difficult to decipher anything meaningful even after the data over the wire is captured. However, it impacts performance slightly and is recommended when the security requirement is very stringent.</p> |
| <i>Color Quality</i> | <p>By default, on a fast Link, the color quality is set to <i>Normal</i> and on a slow link the color quality is set to <i>256 colors</i>. You can change the color quality of the slow link or the fast link to one of the following:</p> <ul style="list-style-type: none"> ◆ 16 Colors: Forces the use of 16-color palette on the managed server during a Remote Management session. This enhances the Remote Management performance particularly over a slow-link. ◆ 256 Colors: Forces the use of 256-color palette on the managed server during a Remote Management session. This enhances the Remote Management performance over a slow-link. ◆ Normal: The color is not altered and the setting is the same on the managed server during a Remote Management session. |
| <i>Network Type</i> | <p>If the managed server is connected by a LAN, select the <i>Fast Links</i> option to enhance the Remote Management performance.</p> <p>If the managed server is connected over a dial-up link or by WAN, select the <i>Slow Links</i> option to enhance the Remote Management performance.</p> |

- 4 To save the *Control Parameter* settings, select the *Save on Exit* check box.

The saved settings are implemented in the next Remote Control session.

Using the Viewing Window Accelerator Keys

You can use accelerator keys to assign shortcut keys to the control options and also to control the display of the Viewing window. Default accelerator key sequences are assigned to each accelerator key option. The Accelerator Keys dialog box displays the default key sequence in the edit field of each accelerator key option. You can define a custom accelerator key sequence to change the default sequence. For more information, see [“Defining a Custom Accelerator Key Sequence” on page 807](#).

To enable the Accelerator Keys option:

- 1 Click the *Remote Management Agent* icon, located at the top left corner of the Viewing window.

- 2 Click *Configure*.
- 3 Select *Enable Accelerator Keys*.

To open the Accelerator Keys dialog box:

- 1 Click the *Remote Management Agent* icon, located at the top left corner of the Viewing window.
- 2 Click *Accelerator Keys*.

Table 20-3 explains the Accelerator Key options you can use to control the display of the Viewing window:

Table 20-3 Accelerator Key Options for Remote Control Session

| Option | Default Keystroke | Description |
|--------------------------------|-------------------|--|
| <i>Toggle Full Screen</i> | Ctrl+Alt+M | Applicable only if the color resolution settings on the management console and managed server are similar. Sizes the Viewing window to the size of your screen without window borders. |
| <i>Refresh Screen</i> | Ctrl+Alt+R | Refreshes the Viewing window. |
| <i>Restart Session</i> | Ctrl+Alt+T | Re-establishes the connection with the managed server. |
| <i>Enable Accelerator Keys</i> | Ctrl+Alt+A | Enables you to change the default accelerator key sequences. |
| <i>Stop Viewing</i> | Left-Shift+Esc | Closes the Viewing window. |
| <i>Configure Dialog</i> | Alt+M | Opens the Control Parameters dialog box. |
| <i>Accelerator Keys Dialog</i> | Alt+A | Opens the Accelerator Keys dialog box. |
| <i>Poll Full Screen</i> | Alt + L | Scans and renders the information of the entire screen. |
| <i>Scale To Fit</i> | Ctrl+Alt+G | Hides the scroll bars and scale the Remote Management window to fit your screen. |
| <i>System Key Pass</i> | Ctrl+Alt+S | Passes Alt-key sequences on the management console to the managed server. |
| <i>Mouse/Keyboard Lock</i> | Ctrl+L | Locks the keyboard and mouse controls at the managed server. This option is available only if the Allow Locking User's Keyboard and Mouse option is enabled in the Server Remote Management Policy . |
| <i>Blank Screen</i> | Ctrl+Alt+B | Blanks the screen at the managed server. This option is available only if the <i>Allow Blanking User's Screen</i> option is enabled in the Server Remote Management Policy . |
| <i>Reboot</i> | Ctrl+Alt+D | Sends the Ctrl+Alt+Del keystroke to the managed server. Display the Security window on the managed server. |

| Option | Default Keystroke | Description |
|----------------------------|-------------------|--|
| <i>Start</i> | Alt+R | Invokes the <i>Start</i> menu on Windows server. |
| <i>Switch Applications</i> | Ctrl+T | Switches applications on managed servers. |

Using the Toolbar Buttons on the Viewing Window

Table 20-4 describes the toolbar options in the Viewing window:

Table 20-4 Viewing Window Toolbar Buttons

| Button | Default Keystroke | Key Function |
|---|-------------------|--|
| <p><i>Screen Blanking</i></p>  | Ctrl+L | <p>Enabled only if the <i>Allow Blanking User's Screen</i> option is enabled in the effective Remote Control policy of the managed server.</p> <p>Blanks the screen at the managed server. When the remote operator selects this option, the screen of the managed server is be blacked out and the operations performed by the remote operator on the managed server are not visible to the user at the managed server.</p> <p>Not supported over certain display adapters. Refer to the ZENworks 7 Server Management Readme (http://www.novell.com/documentation/zenworks7) for the list of display adapters that do not support this feature.</p> |
| <p><i>Mouse and Keyboard Lock</i></p>  | Ctrl+Alt+B | <p>Locks the keyboard and mouse controls at the managed server. When the remote operator selects this option, the user at the managed server will not be able to use the keyboard and mouse controls of the managed server.</p> |
| <p><i>System Start</i></p>  | Alt+R | <p>Invokes the <i>Start</i> menu on the managed server.</p> |
| <p><i>Application Switcher</i></p>  | Ctrl+T | <p>Sends the Alt-tab key sequences to the managed server.</p> <p>Switches applications on managed servers.</p> <p>To switch the applications,</p> <ol style="list-style-type: none"> 1. In the Viewing window, click the <i>Application Switcher</i> icon or press the Application Switcher shortcut key. 2. To traverse to the application use the <i>Application Switcher</i> icon. 3. To view the application, press Tab. |

| Button | Default Keystroke | Key Function |
|--|-------------------|---|
| System Key Pass Through  | Ctrl+Alt+S | Sets the system key pass to On or Off. Passes Alt-key sequences from the management console to the managed server. Certain key sequences such as Ctrl+Esc, Alt+Tab, Ctrl+Alt+Del, and Alt+PrintScreen are not allowed even when the System Key Pass-Through is set to On. However, you can use the toolbar buttons on the Viewing window for the Ctrl+Esc, Alt+Tab, and Ctrl+Alt+Del keystrokes. |
| Reboot  | Ctrl+Alt+D | Sends the Ctrl+Alt+Del keystroke to the managed server. Displays the Security window on the managed server. |
| Refresh Screen  | Ctrl+Alt+R | Refreshes the viewing window. |
| Full Screen Polling  | Alt+L | Scans and renders the information of the entire screen of the managed server continuously. |
| Scale To Fit  | Ctrl+Alt+G | Hides the scroll bars and scales the Remote Management window to fit your screen. |
| Session Encryption  | | Encryption is an optional feature and will be effective per session. If the saved configuration has the option enabled, the session will be encrypted from the start of the session. Encrypting a whole session provides greater security as the data transferred over the wire will be encrypted and it will be difficult to decipher anything meaningful even after the data over the wire is captured. However, it impacts performance slightly and is recommended when the security requirement is very stringent. |

You can define a custom key sequence if you do not want to use the default key sequence. For more information, see [“Defining a Custom Accelerator Key Sequence” on page 807](#).

Enabling the Wallpaper on the Managed Server

When the remote operator initiates a Remote Control session, any wallpaper displayed on the desktop of the managed server is suppressed. This feature reduces the response time from the managed server for requests from the management console because less traffic is generated over the network while the wallpaper is suppressed.

You can configure the control parameter for this option to change the default settings and enable the display of the wallpaper on the managed server. When you terminate the Remote Control session, the suppressed wallpaper will be restored.

To enable the display of suppressed wallpaper on the managed server:

- 1 Click the *Remote Management Agent* icon, located at the top left corner, then click *Configure*.
- 2 Deselect the *Hide Wallpaper* option.

Using the Remote Management Agent Icon

By default, the *Remote Management Agent* icon is displayed in the system tray of the Windows servers. This icon indicates that the Remote Management Agent is loaded on the managed server.

The user at the managed server can right-click the *Remote Management Agent* icon and choose from the following options:

Table 20-5 *Remote Management Icon Options*

| Option | Description |
|--------------------------------|---|
| <i>Terminate RC/RV Session</i> | Disconnects and closes the remote session on the managed server and displays a message on the management console indicating that the remote session is closed. |
| <i>Security</i> | Allows the user at the managed server to set or clear the password for the server. |
| <i>Information</i> | Displays information such as who is accessing the managed server for the remote session, security settings, and the protocol in use for the remote session. For details, see “Obtaining Information About Remote Management Sessions” on page 814 . You can right-click or double-click the <i>Remote Management Agent</i> icon to view the Information window. |
| <i>Shutdown Agent</i> | This option is always dimmed on managed servers. To shut down the Remote Management Agent on managed servers, you must go to the Service Control Panel and stop the “Novell ZENworks Remote Management Agent” service. |
| <i>Request Session</i> | Enables the user at the managed server to request a remote operator to perform remote session. |
| <i>Help</i> | Displays the Remote Management Agent help. |

Setting Up a Password for the Managed Server

The user at the managed server can set an agent password. This password overrides the password set by the administrator during the ZENworks 7 Remote Management installation.

To set the agent password:

- 1 From the managed server, right-click the *Remote Management Agent* icon.
- 2 Click *Security > Set Password*.

Use a password of ten or fewer alphanumeric characters. The password is case sensitive and cannot be blank.

After the completion of the Remote Management session, you can clear the agent password. If you clear the agent password, the remote operator cannot perform the Remote Management operations.

To clear the agent password:

- 1 On the managed server, right-click the *Remote Management Agent* icon.
- 2 Click *Security > Clear Password*.

Obtaining Information About Remote Management Sessions

Using the Information window, the user at the managed server can view details about the session, such as the name of the remote operator how is remotely managing the server, the security settings, and the protocol in use for the remote session.

To view information about remote sessions:

- 1 On the managed server, right-click the *Remote Management Agent* icon.
- 2 Click *Information*.
- 3 Click the *General* tab to view the general information and the *Security* tab to view the security information.

See the following sections for details:

- ◆ “Obtaining General Information” on page 814
- ◆ “Obtaining Security Information” on page 814

Obtaining General Information

Table 20-6 explains the general information you can obtain about Remote Management sessions from the Information window:

Table 20-6 Remote Management Session General Information

| Option | Description |
|--|---|
| <i>RM Operation</i> | Lists the ongoing Remote Management sessions. |
| <i>RM Information > Initiator</i> | Displays the name of the remote operator. |
| <i>RM Information > Protocol</i> | Displays the protocol that the Remote Management Agent uses to communicate with the management console during a remote session. |
| <i>Optimization Status > RC/RV Optimization</i> | Displays if the optimization driver is enabled or disabled for the Remote Management session. The remote session performance is enhanced if the optimization driver is enabled. |

Obtaining Security Information

The Security Information dialog box displays information based on the Remote Control and Remote View sessions.

Table 20-7 Remote Control and Remote View Session Security Information

| Options | Description |
|--------------------------------|--|
| <i>Permission Required</i> | Indicates if the remote operator should obtain permission from the user at the managed server each time the he wants to perform the remote management session on the managed server. |
| <i>Audible Signal Required</i> | Indicates if an audible signal should be sent to the managed server every time the remote operator accesses the managed server. |

| Options | Description |
|-----------------------------------|---|
| <i>Beep Every</i> | Indicates the time interval based on which the audible signal is periodically sent to the managed server. |
| <i>Visual Signal Required</i> | Indicates if a visible signal should be sent to the managed server every time the remote operator accesses the managed server. |
| <i>Session Encryption Enabled</i> | Indicates whether a remote session will be encrypted or not. Session Encryption Enabled is applicable for Remote Control and Remote View. |
| <i>Display Name Every</i> | Indicates the time interval based on which the visual signal is periodically sent to the managed server. |
| <i>Screen Blanking Allowed</i> | Indicates if the remote operator is allowed to blank the managed server screen. Screen Blanking Allowed is applicable for Remote Control only. |
| <i>Locking Control Allowed</i> | Indicates if the remote operator is allowed to lock the keyboard and mouse controls of the managed server. Locking Control Allowed is applicable for Remote Control only. |

20.4.7 Remote Operator Identification Display

The Remote Management Agent will display the identification of the remote operator in the following dialog boxes on the managed server:

- ◆ Permission dialog box
- ◆ Visible signal dialog box

The information displayed is *console_machine_name\console_windows_username*.

20.4.8 Viewing the Audit Log for Remote Management Sessions

ZENworks Server Management records log information on a Windows managed server.

To view the audit log for Remote Management sessions:

- 1 Click *Start > Programs > Administrative Tools > Event Viewer*.
- 2 Click *Log > Application*.
- 3 Double-click the event associated with the source Remote Management Agent.

To view only the events pertinent to the Remote Management Agent, choose *Remote Management Agent* from the *Source* drop-down list in the Filter dialog box.

20.4.9 Improving the Remote Management Performance

The Remote Management performance, especially over a slow link, has been enhanced through using improved compression.

The performance during a Remote Management session over a slow link or a fast link varies depending on the network traffic. For better response time, try one or more of the following strategies:

On the Management Console

- ◆ Select the *Hide Wallpaper* option on the managed server in the Control Parameters dialog box.
- ◆ Assign color settings on the management console higher than the managed server or assign the same color settings for the management console and the managed server.
- ◆ Select *16 Colors* or *256 Colors* mode in the Control Parameters dialog box to enhance the Remote Management performance.
- ◆ The speed of the management console depends upon the processing power of the client machine. We recommended that you to use single-processor client with a Pentium* III, 500MHz (or later).

On the Managed Server

- ◆ Deselect the *Enable Pointer Shadow* option before starting the Remote Control or Remote View session.

To disable Enable Pointer Shadow:

1. From the Windows desktop, click *Start > Settings > Control Panel > double-click Mouse*.
2. Click *Pointers*.
3. Deselect *Enable Pointer Shadow*.
4. Click *Apply > OK*.

- ◆ At the managed server, use a plain background. Do not set a wallpaper pattern.
- ◆ If the Task manager is opened at the target machine, it is recommended to minimize or close it.
- ◆ Make sure that the scrolling texts (such as the debug windows) and animations are not active on the managed server.
- ◆ Make sure to minimize or close the dialog boxes that are not in use.
- ◆ To perform any operations at the managed server, if possible, use the toolbar options instead of menu options.
- ◆ To maximize the Remote Management performance over WAN, configure the following settings in the Control Parameters dialog box at the managed server:
 - ◆ Set the color mode of the managed server to *16 Colors*.
 - ◆ Select the *Slow Link* option.

20.4.10 Shutting Down and Restarting the Remote Management Agent

The following sections explain how you can use the Remote Management Agent during remote sessions:

- ◆ [“Shutting Down the Remote Management Agent” on page 816](#)
- ◆ [“Restarting the Remote Management Agent” on page 817](#)

Shutting Down the Remote Management Agent

You can shut down the Remote Management Agent during a remote session. When you shut down the Remote Management Agent, the remote session stops. To start another remote session, you need

to restart the Remote Management Agent. For more information, see [“Restarting the Remote Management Agent” on page 817](#).

To shut down the Remote Management Agent on a Windows 2000/2003 managed server:

- 1 From the Control Panel, double-click *Administrative Tools*.
- 2 Double-click *Services*.
- 3 Click *Novell ZENworks Remote Management Agent > Stop*.

IMPORTANT: You can stop the Remote Management Agent on Windows 2000/2003 server only if you have the rights to stop the Windows service.

Restarting the Remote Management Agent

During ZENworks Server Management installation, the Remote Management Agent is installed on the managed server and started automatically when the managed server starts up. If you shut down the Remote Management Agent during a remote session, the remote session stops. To start another remote session, you need to restart the Remote Management Agent on the managed server.

To restart the Remote Management Agent on Windows 2000/2003 managed server:

- 1 From the Control Panel, double-click *Administrative Tools*.
- 2 Double-click *Services*.
- 3 Click *Novell ZENworks Remote Management Agent > Start*.

IMPORTANT: You can start the Remote Management Agent on Windows 2000/2003 server only if you have the rights to start the Windows service.

Documentation Updates

P

This section contains information on documentation content changes that have been made in the *Administration* guide for Remote Management since the initial release of Novell® ZENworks® 7 Server Management. The information will help you to keep current on updates to the documentation.

All changes that are noted in this section were also made in the documentation. The documentation is provided on the Web in two formats: HTML and PDF. The HTML and PDF documentation are both kept up-to-date with the documentation changes listed in this section.

The documentation update information is grouped according to the date the changes were published. Within a dated section, the changes are alphabetically listed by the names of the main table of contents sections for Remote Management.

If you need to know whether a copy of the PDF documentation you are using is the most recent, the PDF document contains the date it was published on the front title page or in the Legal Notices section immediately following the title page.

The documentation was updated on the following date:

- ♦ [Section P.1, “July 14, 2006 \(Support Pack 1\),” on page 819](#)
- ♦ [Section P.2, “December 9, 2005,” on page 819](#)

P.1 July 14, 2006 (Support Pack 1)

Updates were made to the following sections. The changes are explained below.

| Location | Change |
|---|------------------------------------|
| Section 20.4.4, “Customizing the Permission Message,” on page 804 | This section has been newly added. |

P.2 December 9, 2005

Page design is reformatted to comply with revised Novell documentation standards.

Management and Monitoring Services

IV

Novell® ZENworks® 7 Server Management provides industry-standards-based monitoring, management, and reporting services for heterogeneous network environments, including support for multi-protocol LAN/WAN networks and servers.

In addition, the Management and Monitoring Services component helps you to pro-actively manage your Novell NetWare®, Windows and Linux servers by responding faster to network problems and increasing overall system availability.

Management and Monitoring Services has the following components:

- ◆ **Novell ConsoleOne®**, which provides the interface where you can manage and administer your network.
- ◆ **Management Site Services**, including:
 - ◆ Alarm Management
 - ◆ Database Administration
 - ◆ MIB Tools Administration
 - ◆ Monitoring Services
 - ◆ Network Discovery
 - ◆ Reporting
 - ◆ Role-Based Services
 - ◆ Topology Mapping
- ◆ **Server Management** for monitoring all the servers in your network.
- ◆ **Traffic Analysis** for monitoring all traffic on Ethernet, token ring, or Fiber Distributed Data Interface (FDDI) network segments.
- ◆ **Advanced Trending Agent** for monitoring.

The Management and Monitoring Services documentation contains the following sections:

- ◆ [Chapter 21, “Configuring Management and Monitoring Services,” on page 823](#)
- ◆ [Chapter 22, “Using Novell ConsoleOne with Management and Monitoring Services,” on page 851](#)
- ◆ [Chapter 23, “Understanding Network Discovery and Atlas Management,” on page 863](#)
- ◆ [Chapter 24, “Understanding Alarm Management,” on page 919](#)
- ◆ [Chapter 25, “Understanding Server Management,” on page 953](#)
- ◆ [Chapter 26, “Using the MIB Tools,” on page 995](#)
- ◆ [Chapter 27, “Using the Probe Manageability Tool,” on page 1019](#)
- ◆ [Chapter 28, “Monitoring Services,” on page 1023](#)
- ◆ [Chapter 29, “Understanding Traffic Analysis,” on page 1029](#)
- ◆ [Chapter 30, “Customizing the Agent Configuration,” on page 1115](#)

- ◆ Chapter 31, “Protocol Decodes Suites Supported by Novell ZENworks Server Management,” on page 1135
- ◆ Chapter 32, “Novell ZENworks Management and Monitoring Services Database,” on page 1143
- ◆ Chapter 33, “Using Reports in Management and Monitoring Services,” on page 1145
- ◆ Chapter 34, “Using SNMP Community Strings,” on page 1159
- ◆ Chapter 35, “Understanding the View Builder,” on page 1165
- ◆ Chapter 36, “Understanding Trap Configuration,” on page 1175
- ◆ Appendix Q, “Setting up Security for Management and Monitoring Services,” on page 1183
- ◆ Appendix R, “Documentation Updates,” on page 1185

Configuring Management and Monitoring Services

21

To use Management and Monitoring Services effectively, you must correctly install and configure the components on your network. You should have already performed a basic installation of Novell® ZENworks® 7 Server Management (see “[Installation](#)” in the *Novell ZENworks 7 Server Management Installation Guide*).

The following sections provide you with the concepts and instructions to help you configure Management and Monitoring Services so that you can use its features to manage your network:

- ♦ [Section 21.1, “Understanding Management and Monitoring Services,”](#) on page 823
- ♦ [Section 21.2, “Planning the Configuration,”](#) on page 828
- ♦ [Section 21.3, “Role-Based Administration,”](#) on page 830
- ♦ [Section 21.4, “Configuring Management and Monitoring Services,”](#) on page 846

21.1 Understanding Management and Monitoring Services

This section provides information about the following components of the Management and Monitoring Services:

- ♦ [Section 21.1.1, “Management Site Services,”](#) on page 823
- ♦ [Section 21.1.2, “Server Management,”](#) on page 826
- ♦ [Section 21.1.3, “Traffic Analysis,”](#) on page 826
- ♦ [Section 21.1.4, “Novell ConsoleOne,”](#) on page 827

21.1.1 Management Site Services

The Management Site Services include the following:

- ♦ [“Network Discovery”](#) on page 824
- ♦ [“Database Administration”](#) on page 824
- ♦ [“Alarm Management”](#) on page 824
- ♦ [“Role-Based Services”](#) on page 824
- ♦ [“Reporting”](#) on page 825
- ♦ [“Topology Mapping”](#) on page 825
- ♦ [“Management Information Base \(MIB\) Tools Administration”](#) on page 826
- ♦ [“Monitoring Services”](#) on page 826

Network Discovery

When network auto discovery is started, the servers, routers, switches which are Simple Network Management Protocol (SNMP) instrumented, and the services hosted on these devices and workstations, are automatically discovered. The discovered data is written to a `.dat` file and displayed in the atlas map on Novell ConsoleOne®.

Maps reflect the scope of discovery set at the management server. By default, all devices that the management server is able to establish communication with, are discovered and stored at the management server. By defining the scope of NetExplorer™, you can limit the number of discovered objects.

For more detailed information on network discovery, see [Chapter 23, “Understanding Network Discovery and Atlas Management,”](#) on page 863.

Database Administration

Novell ZENworks Server Management provides a centralized Common Information Model (CIM)-compliant Sybase* database on the management server. The database serves as a repository for server and network data that can be displayed or formatted in various ways to provide you with the information you need to manage your network. The Server Management data is stored in a topology database containing three logical databases:

- ◆ Topology
- ◆ Alarms
- ◆ Map information

Most database functions are automatic and require very little administration. For more detailed information on Server Management databases, see [Chapter 32, “Novell ZENworks Management and Monitoring Services Database,”](#) on page 1143.

Alarm Management

Alarms recognized by Server Management include SNMP traps, connectivity testing, and threshold profiling. Alarm management processes traps and proprietary alarms and forwards the alarms to Novell ConsoleOne that subscribe to the alarms.

You can perform specific actions on an alarm by specifying the action in the alarm disposition. Some actions, like executing a program, sending an e-mail notification, and creating an archive, audible beep at the Console, and ticker messages, are automatically performed. You can set an action to forward specific processed alarms to other Server Management servers, as well as forward unprocessed SNMP traps directly to a target address of any third-party enterprise management application.

Role-Based Services

Management and Monitoring Services supports role-based administration and task management through Novell eDirectory™. Novell ZENworks Server Management uses Role-based Services to organize Management and Monitoring Services tasks into roles and to assign scope information to a role.

Role-based Services specify tasks that users are authorized to perform. Defining an Role-based Services includes creating an Role-based Services object and specifying the tasks that the role can perform.

For general information on creating Role-based Services objects or specifying tasks that Role-based Services can perform, see [“Configuring Role-Based Administration” on page 844](#).

For information on how Novell ZENworks Server Management implements role-based services, see [Section 21.3, “Role-Based Administration,” on page 830](#).

Reporting

Novell ZENworks Server Management provides reporting services to generate statistical information. These reports can be displayed on Novell ConsoleOne or exported to databases and Web formats. Server Management allows you to generate the following types of reports:

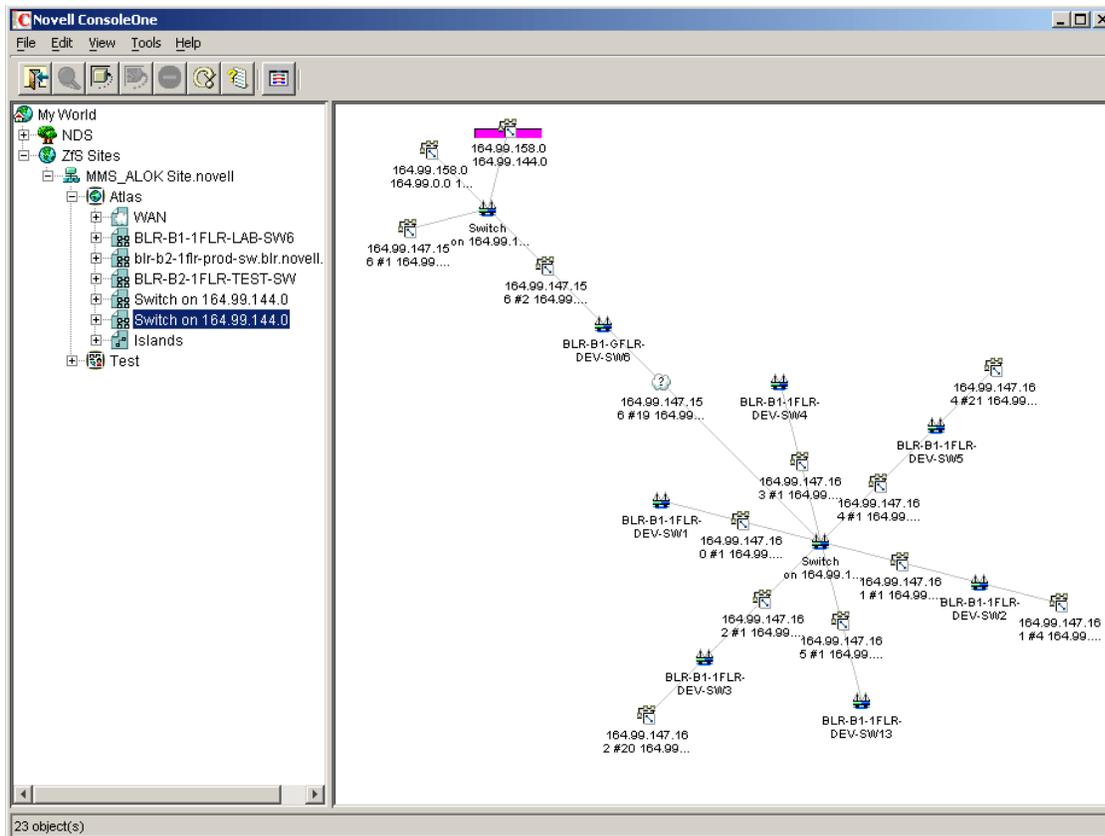
- ◆ Health reports
- ◆ Topology reports
- ◆ Alarm reports

For more detailed information on Management and Monitoring Services reports, see [“Using Reports in Management and Monitoring Services” on page 1145](#).

Topology Mapping

Topology mapping enables you to display maps in the Novell ZENworks Server Management hierarchical atlas as shown in [Figure 21-1](#). Maps reflect the scope of discovery set at the management server.

Figure 21-1 Novell ZENworks Server Management hierarchical atlas



For more detailed information on topology mapping, see [Section 23.3, “Managing the Atlas,”](#) on page 908.

Management Information Base (MIB) Tools Administration

Novell ZENworks Server Management includes the MIB compiler and MIB browser, to manage SNMP devices.

The MIB tools enable you to:

- ◆ Set alarm templates for receiving SNMP traps
- ◆ Display and set values on SNMP devices
- ◆ Update trap definitions in the alarm template database
- ◆ Annotate third-party MIBs

For more detailed information on the MIB tools, see [Chapter 26, “Using the MIB Tools,”](#) on page 995.

Monitoring Services

Monitoring, or SNMP, services include testing the connectivity and availability of a service on a network device. Novell ConsoleOne is notified whenever the status of the service changes. The services that can be monitored include DHCP, DNS, Echo, FTP, HTTP, HTTPS, IP, IPX™, NFS, NNTP, SMTP, SNMP, Time Service, TFTP, and WUser.

For more detailed information on monitoring services, see [Chapter 28, “Monitoring Services,”](#) on page 1023.

21.1.2 Server Management

The server management component enables you to monitor all the servers in your network. This component must be installed on each of the servers you want to monitor using Novell ConsoleOne. During the Novell ZENworks Server Management installation you can select the servers to install the server management component.

You can deploy some or all of the server monitoring software components to meet your management needs best. For more detailed information on server management, see [“Understanding Server Management”](#) on page 953.

21.1.3 Traffic Analysis

The traffic management component provides the traffic analysis services for a NetWare® or Windows* server, to monitor all traffic on an Ethernet, Fiber Distributed Data Interface (FDDI), or token ring network segments.

The traffic analysis services include:

- ◆ Standard and enterprise-specific RFC 1757 MIB descriptions for remote network monitoring
- ◆ Extensions added to Novell eDirectory, including Remote Monitor (RMON) agent configuration
- ◆ Network traffic trending and analysis tools

- ◆ Network health report templates
- ◆ Integration with topology maps
- ◆ Performance threshold configuration and profiling
- ◆ A view of conversations on network segment and utilization
- ◆ Packet capture tools and view

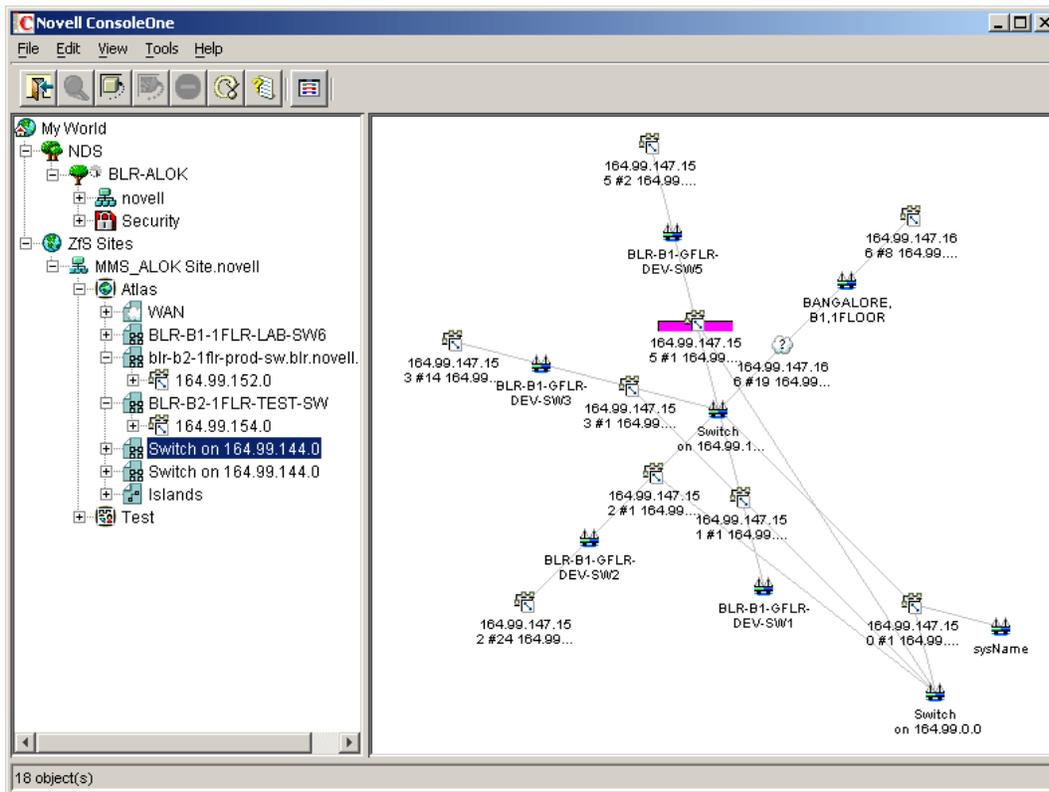
You can deploy some or all of the traffic analysis software components to meet your management needs best. For more detailed information on analyzing the network traffic, see [Chapter 29, “Understanding Traffic Analysis,”](#) on page 1029.

21.1.4 Novell ConsoleOne

The Novell ConsoleOne provides the interface where you can manage and administer your network. Novell ConsoleOne hosts programs (snap-ins) for integrating network administration and management snap-ins, enabling you to manage your network through a single interface.

ZENworks 7 Server Management provides a graphical user interface (GUI) snap-in to the Novell ConsoleOne under the Novell ZENworks Server Management namespace, as shown in [Figure 21-2](#). It provides access to the unique functions provided by Novell ZENworks Server Management.

Figure 21-2 Novell ZENworks Server Management Novell ConsoleOne under the Novell ZENworks Server Management Sites namespace



For more information on Novell ConsoleOne, see the [Novell ConsoleOne Web site \(http://www.novell.com/products/netconsole/consoleone\)](http://www.novell.com/products/netconsole/consoleone).

21.2 Planning the Configuration

This section discusses general planning options for configuring the Management Site Services and some of the Novell ZENworks Server Management agents (alarms, servers, and traffic) on your network. This section also discusses how to plan and implement role-based administration.

Before installing the Management and Monitoring Services software, you must decide what information you need to manage your network effectively. This section contains the following topics to help you decide the kind of information you would need to manage your network.

This section also explains how to configure the Management and Monitoring Services.

- ◆ [Section 21.2.1, “Defining Management Information Needs,” on page 828](#)
- ◆ [Section 21.2.2, “Planning a Strategy to Manage Your Network,” on page 828](#)

This guide also contains specific information on planning server management and segment monitoring in the following sections:

- ◆ [Section 25.2, “Planning for Server Management,” on page 956](#)
- ◆ [Section 29.2, “Planning for Segment Monitoring,” on page 1042](#)

21.2.1 Defining Management Information Needs

Novell ZENworks Server Management is flexible to suit the business needs of different network configurations. You need to understand what information is needed by the groups in your organization and suitably deploy the software to meet those needs.

Typically, the groups in your company may consist of front-line help desk people, back-end information system administrators, and management-level coordinators, who need specific information for planning, budgeting, troubleshooting, and other issues.

For instance, one group might have a set of critical servers that need to be monitored round the clock. You might want real-time monitoring of these servers and receive notification when serious faults occur on these servers. Another example could be a need to generate weekly reports on server trends for a group of defined servers.

21.2.2 Planning a Strategy to Manage Your Network

In order for Novell ZENworks Server Management to monitor and manage devices on your network, it must actively poll your network segments and devices on your network. Novell ZENworks Server Management performs polling of these network objects using standard protocols (SNMP, TCP/IP, and IPX).

The design of the Novell ZENworks Server Management components minimizes the impact on network performance by storing trending information on the servers hosting the SNMP and Remote Monitor (RMON) agents. Polling is directly performed by the management server based on requests coming from connected Novell ConsoleOne.

The Novell ZENworks Server Management system administrator should configure the polling frequency to provide an appropriate level of monitoring for the network environment. A good rule for setting appropriate levels of monitoring is to identify systems that are critical for the operation.

You can then group systems and segments into three basic management categories:

- ♦ **Mission critical:** Segments and devices that need to be actively monitored. Monitoring should be set at a high polling frequency.
- ♦ **Important:** Segments and devices that require less monitoring. These might be systems that host certain services that require a balance between polling overhead and performance. You should set the polling frequency to every few minutes, hours, or days.
- ♦ **Less important:** Segments and devices that require no active monitoring. Polling can be done on-demand to monitor segments and devices, or set to poll infrequently.

Devices that are either not polled or polled infrequently can be configured to send alarms (traps) to the management server to notify errors occurring on the system.

Configuring Your Network

The Management and Monitoring Services components rely on standard network protocols to communicate with devices on your network. In order to discover and accurately monitor your network and its devices, you need to ensure that the communication channels are consistent and well-configured.

The following sections discuss important aspects of your network configuration:

- ♦ [“IP Addressing Strategy” on page 829](#)
- ♦ [“IPX Transport Software” on page 829](#)
- ♦ [“Novell eDirectory and DNS Name Resolution” on page 829](#)
- ♦ [“SNMP Configuration” on page 829](#)

IP Addressing Strategy

If you want to discover devices communicating over IP, ensure that they are configured with a valid IP address to enable you to manage the devices. TCP/IP must be bound on the designated Novell ConsoleOne workstations and IP must be bound on the management server. You can use Dynamic Host Configuration Protocol (DHCP) addressing on Novell ConsoleOne workstation, but a static address must be assigned to the management server.

IPX Transport Software

All devices communicating over IPX that you want to discover and manage must be configured with an IPX/SPX - compatible transport network software stack. NetWare and Windows drivers are included with the operating system installation software. ZENworks Server Management is compatible with the Novell IP Compatibility Mode Driver.

Novell eDirectory and DNS Name Resolution

Verify that your NetWare and Windows servers and network device names are in place before you begin discovering your network. Name resolution can be in the form of local host files, an Novell eDirectory name, or a bindery table. The server names or hostnames are displayed in maps and configuration views rather than in IP or IPX addresses.

SNMP Configuration

The SNMP agents and RMON agents for NetWare and Windows servers and other SNMP-enabled network devices require a community string to be identified on the device. You need to configure

each SNMP-enabled device with a community string and trap target destination that includes that Novell ZENworks Server Management server.

The community strings are used to ensure secure communication between the manager and the agents. In order for the Novell ZENworks Server Management system to communicate with an agent, the community string on the manager and agent must be similar and use the same port. In order to prevent all users from accessing information it is required to change the community string.

If the GET and SET community strings are changed from PUBLIC, you need to change settings at Novell ConsoleOne and on the management server (load NXPCON > SNMP > Add/Edit Community Name) to match the names on your network. For details on how to change the community string, after installing the Management Services, see [“Changing the SNMP Community String” on page 902](#).

For information on configuring the NetWare and Windows server agents, see [Chapter 29, “Understanding Traffic Analysis,” on page 1029](#).

21.3 Role-Based Administration

You can use Novell ConsoleOne, a directory-enabled framework for running Novell network administration utilities. The Novell ZENworks Server Management snap-ins to Novell ConsoleOne fully leverage Novell eDirectory to enable role-based administration and higher levels of security. Through Novell eDirectory, users will be able to log in once and have access to the management components as specified by their roles within their specific scope.

The Server Management snap-ins to Novell ConsoleOne allows you to divide the task of network administration amongst administrators. With Novell ConsoleOne, the functions and tasks of Server Management are organized into different, customized “views” based on each administrator's role in your organization.

The following sections discuss role-based administration:

- ◆ [Section 21.3.1, “Novell ZENworks Management Site,” on page 830](#)
- ◆ [Section 21.3.2, “General Novell ZENworks Server Management Roles,” on page 831](#)
- ◆ [Section 21.3.3, “Novell ZENworks Server Management Role-Based Modules and Roles,” on page 832](#)
- ◆ [Section 21.3.4, “Configuring Role-Based Administration,” on page 844](#)

21.3.1 Novell ZENworks Management Site

The Novell ZENworks management site sets boundaries for accessing object data on the management server through the role-based services. You can create roles and tasks and further define the level of access to network objects and information from the network container space.

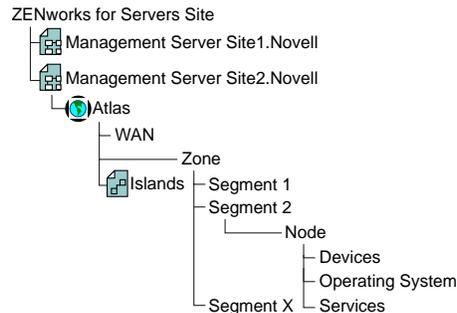
When you install Management and Monitoring Services, a management site, a system administrator role (Role-based Services Admin), and all the site objects are created in Novell eDirectory. A management site defines the scope of objects (networks, segments, routers, bridges, switches, servers, workstations, and so on) discovered on your network. You can create a single site or multiple sites, depending on the size of your network or network management requirements.

A management site could include a single local network configuration or could encompass your entire network. The boundaries of a site are defined by the scope of network discovery. By default,

network discovery is set to discover all connected networks and network nodes. The site object is created in the same context as the server object.

During installation, the default management site that is created is shown in [Figure 21-3](#). A single administration role is established with rights and permissions to all configuration and management tasks in the management system.

Figure 21-3 ZENworks Server Management site



Some default roles that monitor network traffic, handle alarms, and manage server systems, are available and allow you to add users. You can also use them as examples for your new role creations.

In the Server Management role-based services (Role-based Services), permissions that are required to access network objects, configurations, and information are associated with roles. Novell eDirectory User objects can be assigned to appropriate roles. The levels of abstractions in a role are described below:

- ◆ Roles - Created to perform various network management functions in your organization. You can simplify granting of permissions and restrict access to management tools and data by creating appropriate roles.
- ◆ Tasks - Actions performed to utilize components of the management system based on the specific responsibilities.
- ◆ Component/module - A software tool that provides a network management function. Server Management includes components for managing servers, monitoring segment traffic, and providing common services such as database management, alarm handling, and report generation.

The users added to a role, however, retain the access rights, permissions, and policies granted through the Novell eDirectory user account. For example, a user may be granted permission to access and configure a server through Novell eDirectory, but may not be granted permission to manage the server through the Role-based Services in Server Management. Therefore the management role that the user is assigned has limited access to the management services or components/modules in the Novell ZENworks Server Management system.

21.3.2 General Novell ZENworks Server Management Roles

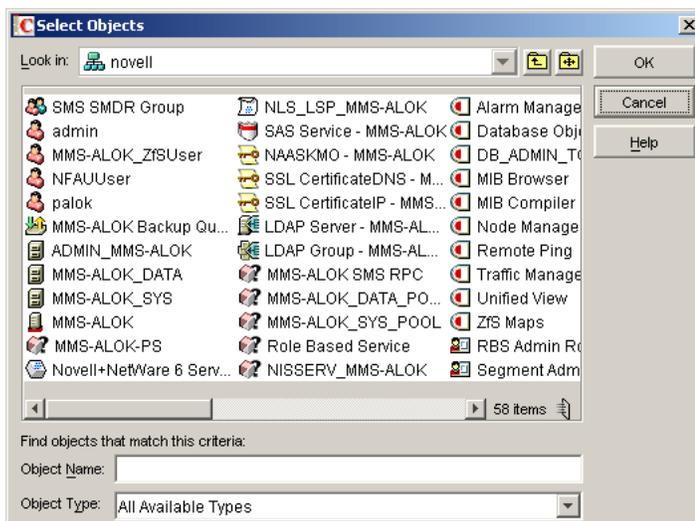
Novell ZENworks Server Management components support role-based services (Role-based Services) and task management through Novell eDirectory. Server Management uses Role-based Services to organize Novell ZENworks Server Management tasks into roles and to assign scope information to a role, user or a group.

Role-based Services roles specify the tasks that users are authorized to perform. Defining an Role-based Services role includes creating an Role-based Services role object and specifying the tasks that the role can perform.

The tasks that Role-based Services roles can perform are displayed as Role-based Services Task objects in your Novell eDirectory tree. These objects are organized into one or more Role-based Services modules, which are containers that correspond to the different Server Management components. As shown in [Figure 21-4](#), Novell ZENworks Server Management provides predefined modules and Role-based Services role objects.

IMPORTANT: You cannot create new modules or tasks. You have to select from the pre-defined modules and tasks that are available.

Figure 21-4 Predefined ZENworks Server Management modules and Role-based Services role objects



You can create any role using the modules and tasks. Each module can have one or more tasks. For example, Role-based Services defines the task for Monitoring Services as Enable Remote Ping. If this task is assigned to your role, you can use the Monitoring Services facility. For a list of the predefined Novell ZENworks Server Management modules and roles along with the associated tasks, see [“Novell ZENworks Server Management Role-Based Modules and Roles” on page 832](#).

For more information on creating role objects using tasks and modules, see [“Configuring Role-Based Administration” on page 844](#).

21.3.3 Novell ZENworks Server Management Role-Based Modules and Roles

This section provides the following tables:

- ◆ [Table 21-1 on page 833](#)
- ◆ [Table 21-2 on page 836](#)

The following table lists each Novell ZENworks Server Management Role-based Services module and the tasks that can be performed for the module.

Table 21-1 ZENworks Server Management Role-based Services module and its associated tasks

| Novell ZENworks Server Management Role-based Services Module | Associated Tasks |
|---|---|
| Alarm Manager | <ul style="list-style-type: none">◆ Add Alarm Note◆ Assign Alarm◆ Define Alarm Disposition◆ Delete Alarm◆ View Active Alarms◆ View Active Alarm History◆ View Alarm Summary |
| Database Object Editor | Database Object Editor |
| DB_Admin_Tool | <ul style="list-style-type: none">◆ DB_BACKUP◆ Database Password Change |
| MIB Browser | Enable MIB Browser |
| MIB Compiler | Enable MIB Compiler |

**Novell ZENworks Server Management
Role-based Services Module****Associated Tasks**

Node Management

- ◆ Clearing a Connection
 - ◆ Create Health Profiles
 - ◆ Create Health Reports
 - ◆ Delete Health Profiles
 - ◆ Delete Health Reports
 - ◆ Downing a Server
 - ◆ Loading an NLM
 - ◆ Mounting and Dismounting a Volume
 - ◆ Read Only All
 - ◆ Read Only All Tabular View
 - ◆ Read Only Health Profiles
 - ◆ Read Only Health Reports
 - ◆ Read Only Homepage
 - ◆ Read Only HostFileSystemView
 - ◆ Read Only InstalledSoftwareView
 - ◆ Read Only Novell NetWareLoadableModuleView
 - ◆ Read Only Novell NetWareUserView
 - ◆ Read Only NetworkPerformanceView
 - ◆ Read Only NTDiskListView
 - ◆ Read Only NTMemoryUsageView
 - ◆ Read Only NTNetworkView
 - ◆ Read Only NTPartitionView
 - ◆ Read Only NTApadpterView
 - ◆ Read Only NTConnectionListView
 - ◆ Read Only NWDiskListView
 - ◆ Read Only NWMemoryUsageView
 - ◆ Read Only NWNetworkMediaView
 - ◆ Read Only NWProtocolView
 - ◆ Read Only NWFileListView
-

**Novell ZENworks Server Management
Role-based Services Module****Associated Tasks**

- ◆ Read Only NWPartitionView
- ◆ Read Only NWQueueJobsListView
- ◆ Read Only NWQueueListView
- ◆ Read Only NWVolumeListView
- ◆ Read Only NWVolumeSegmentView
- ◆ Read Only NWVolumeUsageView
- ◆ Read Only NWRunningSoftwareView
- ◆ Read Only Set Parameter
- ◆ Read Only Trend
- ◆ Read Write All
- ◆ Read Write All TabularView
- ◆ Read Write Health Profiles
- ◆ Read Write Health Reports
- ◆ Read Write Set Parameter
- ◆ Read Write Trend
- ◆ Remote Controlling
- ◆ Restarting a Server
- ◆ Unloading an NLM

Remote Ping**Enable Remote Ping**

Traffic Management

- ◆ Adding_Nodes_For_InactivityMonitoring
 - ◆ Adding_Protocols_For_ProtocalDirectory
 - ◆ Capture_Packets
 - ◆ Deleting_Nodes_For_Inactivity
 - ◆ Deleting_Protocols_For_ProtocolDirectory
 - ◆ Freeing Agent Resources
 - ◆ Setting_Segment_Alarms
 - ◆ View_Conversations
 - ◆ View_LANZ_Agents
 - ◆ View_Protocol_Directory
 - ◆ View_RMON_Summary
-

| Novell ZENworks Server Management Role-based Services Module | Associated Tasks |
|---|---|
| | <ul style="list-style-type: none"> ◆ View_Segment_Alarms ◆ View_Segment_Dashboard ◆ View_Segment_Monitor_Nodes_For_Inactivity ◆ View_Segment_Protocol_Distribution ◆ View_Segment_Stations ◆ View_Segment_Summary ◆ View_Segment_Trends ◆ View_Switch_Port_Traffic ◆ View_Switch_Summary |
| Unified View | <ul style="list-style-type: none"> ◆ Unified View for Devices ◆ Unified View for Segments |
| Novell ZENworks Server Management Maps | <ul style="list-style-type: none"> ◆ Import ◆ Layout ◆ Print ◆ Rebuild ◆ Rename ◆ Save |

The following table lists each predefined Novell ZENworks Server Management Role-based Services and the specific tasks that can be performed for each of the roles:

Table 21-2 *Predefined ZENworks Server Management Role-based Services and Modules*

| Management and Monitoring Services Predefined Role-based Services Role | Management and Monitoring Services Role-based Services Module | Assigned Default Tasks |
|---|--|-------------------------------|
| Role-based Services_Administrator | All Modules | All available tasks |

| Management and Monitoring Services Predefined Role-based Services Role | Management and Monitoring Services Role-based Services Module | Assigned Default Tasks |
|--|---|---|
| Segment_Administrator | Alarm Manager | <ul style="list-style-type: none"> ◆ View Alarm Summary ◆ View Active Alarms ◆ View Alarm History ◆ Assign Alarms ◆ Add Alarm Note |
| | DM_Admin_Tool | No available tasks |
| | MIB Browser | No available tasks |
| | MIB Compiler | Enable MIB Compiler |
| | Node Management | <ul style="list-style-type: none"> ◆ Read Only Health Profiles ◆ Read Only Health Reports |
| | Remote Ping | Enable Remote Ping |
| | Traffic Management | <ul style="list-style-type: none"> ◆ Adding_Nodes_For_InactivityMonitoring ◆ Adding_Protocols_For_ProtocolDirectory ◆ Capture_Packets ◆ Setting_Segment_Alarms ◆ View_Conversations ◆ View_LANZ_Agents ◆ View_Protocol_Directory ◆ View_RMON_Summary ◆ View_Segment_Alarms ◆ View_Segment_Dashboard ◆ View_Segment_Monitor_Nodes_For_Inactivity ◆ View_Segment_Protocol_Distribution ◆ View_Segment_Stations ◆ View_Segment_Summary ◆ View_Segment_Trends ◆ View_Switch_Port_Traffic ◆ View_Switch_Summary |
| | Novell ZENworks Server Management Maps | <ul style="list-style-type: none"> ◆ Layout ◆ Print |
| | Unified Views | Unified Views for Segments |

| Management and Monitoring Services Predefined Role-based Services Role | Management and Monitoring Services Role-based Services Module | Assigned Default Tasks |
|--|---|--|
| Segment Manager | Alarm Manager | <ul style="list-style-type: none"> ◆ Assign Alarms ◆ Define Alarms Disposition ◆ Delete Alarms ◆ View Alarm Summary ◆ View Active Alarms ◆ View Alarm History ◆ Add Alarm Note |
| | DM_Admin_Tool | No available tasks |
| | Database Object Editor | Database Object Editor |
| | MIB Browser | Enable MIB Browser |
| | MIB Compiler | Enable MIB Compiler |
| | Node Management | <ul style="list-style-type: none"> ◆ Create Health Profiles ◆ Create Health Reports ◆ Delete Health Profiles ◆ Delete Health Reports ◆ Read Write Health Profiles ◆ Read Only Health Profiles ◆ Read Write Health Reports ◆ Read Only Health Reports |
| | Remote Ping | Enable Remote Ping |

| Management and Monitoring Services Predefined Role-based Services Role | Management and Monitoring Services Role-based Services Module | Assigned Default Tasks |
|--|---|---|
| Segment Manager <i>continued</i> | Traffic Management | <ul style="list-style-type: none"> ◆ Adding_Nodes_For_InactivityMonitoring ◆ Adding_Protocols_For_ProtocolDirectory ◆ Capture_Packets ◆ Deleting_Nodes_For_InactivityMonitoring ◆ Deleting_Protocols_For_ProtocolDirectory ◆ Freeing Agent Resources ◆ Setting_Segment_Alarms ◆ View_Conversations ◆ View_LANZ_Agents ◆ View_Protocol_Directory ◆ View_RMON_Summary ◆ View_Segment_Alarms ◆ View_Segment_Dashboard ◆ View_Segment_Monitor_Nodes_For_Inactivity ◆ View_Segment_Protocol_Distribution ◆ View_Segment_Stations ◆ View_Segment_Summary ◆ View_Segment_Trends ◆ View_Switch_Port_Traffic ◆ View_Switch_Summary |
| | Novell ZENworks Server Management Maps | <ul style="list-style-type: none"> ◆ Import ◆ Layout ◆ Print ◆ Rebuild ◆ Rename ◆ Save |

| Management and Monitoring Services Predefined Role-based Services Role | Management and Monitoring Services Role-based Services Module | Assigned Default Tasks |
|--|---|---|
| Segment Monitor | Alarm Manager | <ul style="list-style-type: none"> ◆ View Alarm Summary ◆ View Active Alarms ◆ View Alarm History |
| | DM_Admin_Tool | No available tasks |
| | MIB Compiler | No available tasks |
| | MIB Browser | No available tasks |
| | Node Management | <ul style="list-style-type: none"> ◆ Read Only Health Profiles ◆ Read Only Health Reports |
| | Remote Ping | Enable Remote Ping |
| | Traffic Management | <ul style="list-style-type: none"> ◆ Capture_Packets ◆ View_Conversations ◆ View_LANZ_Agents ◆ View_Protocol_Directory ◆ View_RMON_Summary ◆ View_Segment_Alarms ◆ View_Segment_Dashboard ◆ View_Segment_Monitor_Nodes_For_In activity ◆ View_Segment_Protocol_Distribution ◆ View_Segment_Stations ◆ View_Segment_Summary ◆ View_Segment_Trends ◆ View_Switch_Port_Traffic ◆ View_Switch_Summary |
| | Novell ZENworks Server Management Maps | <ul style="list-style-type: none"> ◆ Layout ◆ Print |
| | Unified Views | Unified View for Segments |

| Management and Monitoring Services Predefined Role-based Services Role | Management and Monitoring Services Role-based Services Module | Assigned Default Tasks |
|--|---|---|
| Server Administrator | Alarm Manager | <ul style="list-style-type: none"> ◆ Assign Alarm ◆ Define Alarm Disposition ◆ Delete Alarm ◆ View Alarm Summary ◆ View Active Alarms ◆ View Alarm History ◆ Add Alarm Note |
| | DM_Admin_Tool | No available tasks |
| | MIB Browser | Enable MIB Browser |
| | MIB Compiler | No available tasks |
| | Node Management | <ul style="list-style-type: none"> ◆ Clearing a Connection ◆ Loading an NLM ◆ Mounting and Dismounting a Server Volume ◆ Downing a Server ◆ Read Only Health Profiles ◆ Read Only Health Reports ◆ Read Write All ◆ Restarting a Server ◆ Unloading an NLM |
| | Remote Ping | Enable Remote Ping |
| | Traffic Management | No available tasks |
| | Novell ZENworks Server Management Maps | <ul style="list-style-type: none"> ◆ Layout ◆ Print |
| | Unified Views | Unified Views for Devices |

| Management and Monitoring Services Predefined Role-based Services Role | Management and Monitoring Services Role-based Services Module | Assigned Default Tasks | |
|--|---|--|--------------------------|
| Server Manager | Alarm Manager | <ul style="list-style-type: none"> ◆ Assign Alarm ◆ Define Alarm Disposition ◆ Delete Alarm ◆ View Alarm Summary ◆ View Active Alarms ◆ View Alarm History ◆ Add Alarm Note | |
| | DM_Admin_Tool | No available tasks | |
| | MIB Browser | No available tasks | |
| | MIB Compiler | No available tasks | |
| | Node Management | <ul style="list-style-type: none"> ◆ Clearing a Connection ◆ Create Health Profiles ◆ Create Health Reports ◆ Delete Health Profiles ◆ Delete Health Reports ◆ Downing a Server ◆ Loading an NLM ◆ Mounting and Dismounting a Server Volume ◆ Read Only Health Profiles ◆ Read Only Health Reports ◆ Read Write All ◆ Read Write Health Profiles ◆ Read Write Health Reports ◆ Restarting a Server ◆ Unloading an NLM | |
| | Remote Ping | No available tasks | |
| | Traffic Management | No available tasks | |
| | Novell ZENworks Server Management Maps | <ul style="list-style-type: none"> ◆ Import ◆ Layout ◆ Print ◆ Rebuild ◆ Rename ◆ Save | |
| | Server Manager continued | Database Object Editor | Database Object Editor |
| | | Unified Views | Unified View for Devices |

| Management and Monitoring Services Predefined Role-based Services Role | Management and Monitoring Services Role-based Services Module | Assigned Default Tasks |
|--|---|---|
| Server Monitor | Alarm Manager | <ul style="list-style-type: none"> ◆ View Alarm Summary ◆ View Active Alarms ◆ View Alarm History |
| | DM_Admin_Tool | No available tasks |
| | MIB Browser | No available tasks |
| | MIB Compiler | No available tasks |
| | Node Management | <ul style="list-style-type: none"> ◆ Read Only Health Profiles ◆ Read Only Health Reports ◆ Read Only Homepage ◆ Read Only HostFileSystemView ◆ Read Only InstalledSoftwareView ◆ Read Only Novell NetWareLoadableModulesView ◆ Read Only Novell NetWareUserView ◆ Read Only NetworkPerformanceView ◆ Read Only NTDiskListView ◆ Read Only NTMemoryUsageView ◆ Read Only NTNetworkView ◆ Read Only NWConnectionListView ◆ Read Only NWOpenListView ◆ Read Only NWDiskListView ◆ Read Only NWMemoryUsageView ◆ Read Only NWNetworkMediaView ◆ Read Only NWFileListView ◆ Read Only NWWVolumeListView ◆ Read Only NWWVolumeUsageView ◆ Read Only RunningSoftwareView ◆ Read Only Trend |
| | Remote Ping | Enable Remote Ping |
| | Traffic Management | No available tasks |
| | Novell ZENworks Server Management Maps | <ul style="list-style-type: none"> ◆ Layout ◆ Print |

| Management and Monitoring Services Predefined Role-based Services Role | Management and Monitoring Services Role-based Services Module | Assigned Default Tasks |
|--|---|---|
| Site Database Administrator | Alarm Manager | No available tasks |
| | DM_Admin_Tool | <ul style="list-style-type: none"> ◆ DB_BACKUP ◆ Database Password Change |
| | MIB Browser | No available tasks |
| | MIB Compiler | No available tasks |
| | Node Management | No available tasks |
| | Remote Ping | No available tasks |
| | Traffic Management | No available tasks |
| | Novell ZENworks Server Management Maps | No available tasks |

21.3.4 Configuring Role-Based Administration

Defining an Role-based Services role includes creating an Role-based Services role object and specifying the tasks that the role can perform.

The following sections discuss how to configure Role- Based Administration:

- ◆ “Defining Role-based Services Role” on page 844
- ◆ “Creating an External Scope” on page 845
- ◆ “Configuring a Scope Object” on page 845
- ◆ “Assigning Role-based Services Role Membership and Scope” on page 845

Defining Role-based Services Role

Role-based Services roles specify the tasks that users are authorized to perform in specific administration applications. Defining an Role-based Services role includes the following sections:

- ◆ “Creating an Role-based Services Role Object” on page 844
- ◆ “Specifying the Tasks that Role-based Services Roles Can Perform” on page 845

Creating an Role-based Services Role Object

To create an Role-based Services role object:

- 1** In Novell ConsoleOne, right-click the container that you want to create the Role-based Services role object, then click *New > Object*.
- 2** In *Class*, select *Role-based Services:Role*, then click *OK*.
- 3** Enter a name for the new Role-based Services role object.

Ensure to follow proper Novell eDirectory naming conventions. For Novell eDirectory naming conventions see [Novell eDirectory Administration Guide \(http://novell.com/documentation\)](http://novell.com/documentation).

Example: Password Administrator Role.

- 4 Click *OK*.

Specifying the Tasks that Role-based Services Roles Can Perform

To specify the tasks:

- 1 In Novell ConsoleOne, right-click an Role-based Services role, then click *Properties*.
Role-based Services task objects are located only in Role-based Services module containers
- 2 In the *Role Based Services* tab, make the associations you want.
- 3 Select the *Role Content* page, then add the list of tasks that the role can perform.
- 4 Click *OK*.

Creating an External Scope

To create an external scope:

- 1 In Novell ConsoleOne, right-click the container that you want to create the scope object, then click *New > Object*.
- 2 In *Class*, select *MW:Scope*, then click *OK*.
- 3 Enter a name for the new *MW:Scope* object.
Ensure to follow proper Novell eDirectory naming conventions. For Novell eDirectory naming conventions see [Novell eDirectory Administration Guide \(http://novell.com/documentation\)](http://novell.com/documentation).
Example: Password Administrator Role.
- 4 Click *OK*.

Configuring a Scope Object

To configure a scope object:

- 1 In Novell ConsoleOne, right-click the scope object, then click *Properties*.
- 2 Browse the site object to which the scope is associated.
- 3 In the Site scope, browse to select the computers to the site scope.
- 4 In the SQL script, specify the scope by selecting the object and the operator from the drop-down list.
- 5 Click *OK*.

IMPORTANT: By default the scope object will have all-site access.

The effective scope will be a union of Site scope and the objects specified in SQL script.

Assigning Role-based Services Role Membership and Scope

To assign an Role-based Services role and scope to a user:

- 1 In Novell ConsoleOne, right-click the user object to which you want to assign the role and scope, then click *Properties*.
- 2 Click *Role Based Services* tab, then click *Assigned Roles*.
- 3 Click *Add* to add the required role to the user.

- 4 Click *Scope* to add the scope for the user.
- 5 Click *OK*.

IMPORTANT: If a user is assigned two different roles with different scopes, the user has rights to all the tasks (union of tasks in role1 and tasks in role2) irrespective of the scopes.

You cannot assign role and scope to User groups and Organization Unit.

21.4 Configuring Management and Monitoring Services

Novell ZENworks Server Management is made up of several components, some of which require certain setup tasks before you can use them, and others that do not.

The following components do not require any specific setup tasks:

- ◆ Novell ZENworks Server Management databases
- ◆ Role-based services (Role-based Services)
- ◆ Management Information Base (MIB) tools
- ◆ Novell ConsoleOne
- ◆ Reporting
- ◆ SNMP services

The following sections describe the setup tasks that are required to get the following components up and running:

- ◆ [Section 21.4.1, “Stopping and Starting Management and Monitoring Services,” on page 846](#)
- ◆ [Section 21.4.2, “Setting Up Discovery and Starting Back-End Processes,” on page 847](#)
- ◆ [Section 21.4.3, “Setting Up the Alarm Management System,” on page 848](#)
- ◆ [Section 21.4.4, “Setting Up Monitoring,” on page 848](#)
- ◆ [Section 21.4.5, “Setting Up the Traffic Analysis Agent,” on page 848](#)

21.4.1 Stopping and Starting Management and Monitoring Services

If you need to install other software or perform other maintenance functions on your server, you can stop Management and Monitoring Services and down the server. After performing the maintenance, you must reboot the server and restart the services in order for the server to resume its Management and Monitoring Services.

To stop and start Management and Monitoring Services, complete the following steps at the management server console prompt:

- 1 To stop and unload Management and Monitoring Services, enter `stopmms`.
- 2 To stop all JAVA processes, enter `java -kill`.
- 3 To exit JAVA, enter `java -exit`.
- 4 To restart the server, enter `restart server`.

To down the server, enter `down server`. You need to start the server again.

- 5 To stop and unload all management and monitoring services and naming service, enter `stopmms -n`

Because the appropriate commands to start the back-end and discovery processes (SLOADER and NETEXPLOR) were inserted in the `autoexec.ncf` file when you installed Management and Monitoring Services, restarting the server will start these processes. If you modified the `autoexec.ncf` file and need to manually start these processes, see [“Manually Starting Discovery and Back-End Processes” on page 848](#).

21.4.2 Setting Up Discovery and Starting Back-End Processes

The discovery software on the management server automatically discovers the nodes on your network. Network nodes include servers, desktops, routers, switches, and any other network devices. Discovery starts automatically when the Novell ZENworks Server Management software is loaded on the management server and runs continually, 24 hours a day. The amount of time to build a complete database depends on the size of your network. Very small networks might take one or two hours; very large networks (several thousand nodes) might require several days.

It is recommended that you run Network discovery on a standalone as the discovery process consumes a longer duration if you use the system.

After installation, your servers are in one of the following states:

- ◆ Discovery and back-end services are running.

If you choose to start the auto discovery process and back-end services during installation, discovery is running on your Novell ZENworks Server Management server and your network is continually being discovered. You do not need to do anything further with regards to configuring discovery unless you want to modify your discovery parameters after you check the results of the initial discovery. For instructions on checking the results of discovery and modifying your discovery parameters, see [Chapter 23, “Understanding Network Discovery and Atlas Management,” on page 863](#).

IMPORTANT: After modifying any discovery parameters, you must restart the server as described in [“Stopping and Starting Management and Monitoring Services” on page 846](#).

- ◆ Discovery and back-end services are not running.

If you selected No, and did not start the auto discovery process and back-end services during installation, you must start discovery after you modify the default discovery parameters. For specific instructions on modifying discovery parameters, see [Chapter 23, “Understanding Network Discovery and Atlas Management,” on page 863](#).

Before discovering your network, you can modify the following discovery parameters:

- ◆ SNMP Community Strings. Ensure that discovery is configured with the community strings of your devices.
- ◆ Discovery Scope. By default, discovery will discover the entire network if correct community strings are provided. If the discovery scope needs to be limited for some reason, it can be modified.
- ◆ IPX Discovery. IPX discovery will take place as long as the Novell ZENworks Server Management server has a valid IPX address binding. If there is no IPX address bound to

the Server Management server, but there are IPX networks that need to be discovered, install the NetWare server in CMD mode (load SCMD).

IMPORTANT: After modifying any discovery parameters, you must restart the services as described in [“Stopping and Starting Management and Monitoring Services” on page 846](#). If you never started discovery or the back-end services, you can manually start the services as described in [“Manually Starting Discovery and Back-End Processes” on page 848](#).

Manually Starting Discovery and Back-End Processes

The commands to start auto discovery and load the back-end services are inserted into the `autoexec.ncf` file by the installation program. Restarting the server will automatically start these processes. However, if you remove these commands you will need to manually start auto discovery and load the back-end services (management site services).

During installation, a search path is added to the `autoexec.ncf` file to the management server program file path — `Novell ZENworks\mms\mwserver\bin`

Enter `startmms.ncf` at the server to start the discovery and backend process.

The server will accept requests from Novell ConsoleOne only after the backend processes have been completely loaded.

21.4.3 Setting Up the Alarm Management System

The Novell ZENworks Server Management Alarm Management System can receive SNMP traps from any SNMP-enabled device or computer hosting a proxy SNMP agent. If your network device is using Management Agent for NetWare, Management Agent for Windows, Traffic Analysis Agent for NetWare, or Traffic Analysis Agent for Windows software, the device is discovered automatically for you. No setup is needed after installing the software.

Third-party SNMP agents require some setup before traps can be received. For information on setting up third-party SNMP agents, see [“SNMP Configuration” on page 829](#).

21.4.4 Setting Up Monitoring

Because the Management Agent for NetWare and the ManageWise[®] Agent for Windows are based on SNMP, all actions that are directed from network management console to a server involve SNMP SET and GET requests from the manager to the agent. Any Novell ConsoleOne requesting data from a managed server does so by issuing an SNMP GET request. An SNMP SET command is required to set server alarm thresholds or configuration parameters. Conducting these management operations from Novell ConsoleOne such as Novell ConsoleOne, raises the issue of ensuring security. In particular, unauthorized users setting configuration parameters on a server could cause performance problems or even sabotage network operations.

For these reasons, you should secure communication between the management system and your SNMP agents. For further information on SNMP security, [“SNMP Configuration” on page 829](#).

21.4.5 Setting Up the Traffic Analysis Agent

The Traffic Analysis Agent for NetWare is a distributed network analyzer that complements Novell ZENworks Server Management. While other Novell ZENworks Server Management agents collect

data about specific network nodes, such as servers, the Traffic Analysis Agent for NetWare observes the interaction among these nodes on a specific LAN segment.

The agent is installed on a NetWare server. To set up Traffic Analysis Agent for NetWare, see [“Starting the Traffic Analysis Agent for NetWare” on page 849](#).

The Traffic Analysis Agent for Windows uses SNMP to communicate with the management server. After installation, in order for the Traffic Analysis Agent for Windows to operate, you must start the SNMP services. To start SNMP services, complete [“Starting the SNMP Service for the Traffic Analysis Agent for Windows” on page 849](#).

After the agents are set up, you must restart the Windows server on which the agent resides.

Starting the Traffic Analysis Agent for NetWare

The installation program for the Traffic Analysis Agent for NetWare modifies the `autoexec.ncf` file so that the agent starts automatically. Therefore, you do not need any further configuration. If, however, you are upgrading from a previous version of the Traffic Analysis Agent (referred to as the Traffic Analysis agent), and did not uninstall the previous version, you must ensure that each server on which you upgraded the agent will run the new Traffic Analysis Agent.

To ensure that the upgraded NetWare servers run the new Traffic Analysis Agent:

- 1 On each NetWare server where you upgraded the Novell ZENworks Server Management Traffic Analysis Agent, open the `autoexec.ncf` file located in `sys:\system`.
- 2 Comment out the following lines by placing a `#` character at the beginning of the line as follows:

```
#Search add lanzdir
#LANZ.NCF
```

The first statement defines the search path where *lanzdir* is the directory in which the older agent is installed. The second statement loads the older agent.

- 3 Save the file and restart the server.

The new agent will load and run automatically. The `lanz.ncf` file in the `agentinstallfolder\lanz` will start the Traffic Analysis agent. The `ulanz.ncf` in the same folder will stop the Traffic Analysis agent.

Starting the SNMP Service for the Traffic Analysis Agent for Windows

If you have configured Windows to start the SNMP service automatically, the agent installed on Windows starts with the SNMP service when you start Windows.

If you have not configured Windows to start the SNMP service automatically, do either of the following:

- ◆ At the command prompt, enter `net start snmp`.
- ◆ On Windows 2000/2003: In the Windows Control Panel, double-click *Administrative Tools* > *Services*, select *SNMP* from the list of services, then click *Start*.

When the SNMP service is started, the Traffic Analysis Agent for Windows will also start.

Using Novell ConsoleOne with Management and Monitoring Services

22

The Novell® ZENworks® Server Management console is a snap-in to the Novell ConsoleOne® management tool. Novell ZENworks Server Management expands Novell ConsoleOne management capabilities by adding menu options, property pages for existing Novell eDirectory™ objects, and ways to browse and organize network resources. This section introduces Novell ConsoleOne features that are unique to Server Management, including:

- ♦ Section 22.1, “Navigating the Novell ZENworks Server Management Namespace,” on page 851
- ♦ Section 22.2, “Selecting Novell ZENworks Server Management Options,” on page 853
- ♦ Section 22.3, “Working with Views,” on page 854

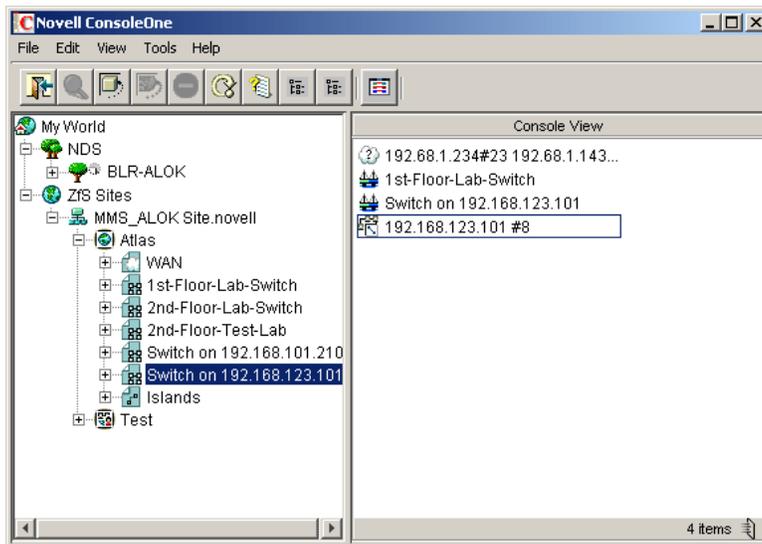
For more information on basic Novell ConsoleOne capabilities, see the [Novell ConsoleOne Administration Guide](http://novell.com/documentation) (<http://novell.com/documentation>).

22.1 Navigating the Novell ZENworks Server Management Namespace

In Novell ConsoleOne, your network and its resources are regarded as a set of objects and are arranged in various containers. Each top-level object is referred to as a namespace. To view your network and its resources on Novell ConsoleOne, you must log in to the Novell eDirectory tree which contains Management Site Server object.

The Novell ZENworks Server Management Novell ConsoleOne snaps in to Novell ConsoleOne under the Novell ZENworks Server Management Sites namespace, as shown in [Figure 22-1](#):

Figure 22-1 ZENworks Server Management ConsoleOne snap-ins



In general, you can perform administration tasks by browsing to an object in the left frame, right-clicking it, and clicking an option. Objects within the Novell ZENworks Server Management namespace are arranged in the following hierarchy:

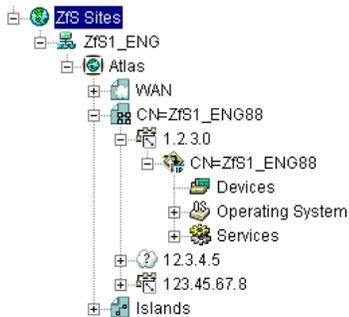
1. **Novell ZENworks Server Management sites object:** This is the Novell ZENworks Server Management namespace container. It is the top of the Novell ZENworks Server Management namespace hierarchy. Expand this object to display a list of Server Management management sites.
2. **Management Site:** This object represents a Novell ZENworks management server. It represents an Novell eDirectory object that defines a collection of discovered objects that collectively make up a group of services. Expand this object to display the atlas for the site.
3. **Atlas:** This is the container object for all discovered topology objects. The atlas can contain the following types of pages:
 - ◆ **WAN page:** Summarizes the entire network.
 - ◆ **Area page:** Displays segments on the network. There may be more than one Area page, depending on how your network is organized.
 - ◆ **Islands page:** Displays segments with undetermined connectivity.
4. **Segments:** Within each atlas page is a listing of the segment objects that are included in that section of the atlas.
5. **Nodes:** Within each segment object is a listing of server and node objects that reside on the segment. The icon displayed varies by the node type.
6. **Node Details:** Expand a node object to display a list of system internal components. Server data is grouped into the following three categories:
 - ◆ **Devices**
 - ◆ **Operating System**
 - ◆ **Services**

You can drill down into the server configuration further by clicking the plus signs next to the Devices, Operating System, and Services objects to display details about the internal

components of the server. The internal components include the processors, installed software, volumes, kernel, and adapters associated with the server. For more details about the node objects, see [“Object Hierarchy” on page 974](#).

Figure 22-2 illustrates the Novell ZENworks Server Management namespace hierarchy:

Figure 22-2 ZENworks Server Management namespace hierarchy



22.2 Selecting Novell ZENworks Server Management Options

To display the Novell ZENworks Server Management options, you want to monitor or manage in the left frame, right-click the object. The options available are displayed. Novell ZENworks Server Management provides three main options:

- ◆ [Section 22.2.1, “Views,” on page 853](#)
- ◆ [Section 22.2.2, “Properties,” on page 854](#)
- ◆ [Section 22.2.3, “Actions,” on page 854](#)

22.2.1 Views

Views are different ways of displaying information. Novell ZENworks Server Management provides a variety of views designed to help you view the information of your network in different ways. The views Server Management provides are:

- ◆ **Atlas**: Provides a graphical representation of the discovered network topology, the physical location of nodes, node configuration, and alarm information.
- ◆ **Console**: Displays the objects contained in the selected container object. This view is useful while navigating the Novell ZENworks Server Management site.
- ◆ **Trend**: Provides a graphical representation of current and historical trend data by hour, day, week, month, or year. Monitoring trend data helps you with tasks such as determining which server is being used, who is using the server, troubleshooting problems, balancing load across multiple servers, and planning resources.
- ◆ **Active Alarms**: Provides a tabular display of alarm statistics for all the current alarms received from segments or devices, per management site. This view is refreshed whenever a new alarm occurs on the network.
- ◆ **Alarm History**: Provides a tabular display of all archived alarms, including the handled status of each alarm. This view is refreshed whenever a new alarm occurs on the network.

- ◆ **Alarm Summary:** Provides a graphical representation of the summary of alarms you have received. The view is divided into three panels of representation: pie chart panel, bar graph pane, and trend panel. Provides a tabular display of all archived alarms, including the handled status of each alarm.
- ◆ **Summary:** Provides a tabular information about the selected object's configuration. For example, the summary view for a server object displays information about NetWare Loadable Module™ files, memory usage, adapters, network interfaces, disks and disk controllers, volumes, queues, users, connections, open files, alarms, and installed software.

In addition to these main views, Novell ZENworks Server Management provides additional views for many of the objects in the hierarchy. For example, if you select a memory object, you can select a disk cache view that displays utilization for disk cache memory. For more information on the available views and the specific information displayed in an object view, see [“Object View Details” on page 975](#).

22.2.2 Properties

The Novell ZENworks Server Management Novell ConsoleOne provides several property pages that allow you to control Novell ZENworks Server Management-specific settings. To access the Novell ZENworks Server Management property pages, right-click an object and then click Properties.

- ◆ At the site level, Novell ZENworks Server Management provides property pages that allow you to edit global properties like Alarm Dispositions, Novell ZENworks Server Management Database settings, Simple Network Management Protocol (SNMP) settings, Management Information Base (MIB) Pool entries, and health report profiles.
- ◆ At the server level, Novell ZENworks Server Management provides property pages that allow you to modify SNMP settings.

For general information on using Novell ConsoleOne property pages, see the [Novell ConsoleOne Administration Guide \(http://novell.com/documentation\)](http://novell.com/documentation).

22.2.3 Actions

You can perform one or more actions on some objects. For example, if you right-click a server object, the *Actions* menu provides options for restarting or shutting down the server. However, if you right-click a volume object, the *Actions* menu provides options for mounting or dismounting the volume. For more information on performing actions on a managed object, see [“Executing Server Commands” on page 972](#).

22.3 Working with Views

Novell ZENworks Server Management Novell ConsoleOne provides two types of views: tabular (list) views and graphical views. The Console, Active Alarms, and Alarm History views are all tabular views. The atlas and Trend views are both graphical views. The Summary view may contain both tabular and graphical elements.

There are many characteristics that are common to all views. This section describes the common tasks you can perform on the Server Management views, including:

- ◆ [Section 22.3.1, “Changing the Appearance of a View,” on page 855](#)

- ◆ Section 22.3.2, “Modifying Columns,” on page 856
- ◆ Section 22.3.3, “Filtering Views,” on page 857
- ◆ Section 22.3.4, “Sorting Views,” on page 858
- ◆ Section 22.3.5, “Printing a View,” on page 858
- ◆ Section 22.3.6, “Exporting a View,” on page 858
- ◆ Section 22.3.7, “Saving Views,” on page 859
- ◆ Section 22.3.8, “Deleting and Renaming Custom Views,” on page 859
- ◆ Section 22.3.9, “Displaying Multiple Views in Novell ConsoleOne Views,” on page 860

22.3.1 Changing the Appearance of a View

In a view, you can change the following:

- ◆ “Changing the Display Font” on page 855
- ◆ “Customizing Grid Lines” on page 855
- ◆ “Displaying the View Title” on page 856

Changing the Display Font

To change the font of the text on a tabular view's headings or rows:

- 1 Click *View > Settings > Appearance*.

The Appearance dialog box is displayed.

- 2 To change the header or row font, click the appropriate button as follows:
 - ◆ To change the header font, click the *Header Font* button.
 - ◆ To change the row font, click the *Row Font* button.

The Fonts dialog box is displayed.

- 3 Select the font options you want, then click *OK* to close the Fonts dialog box.
- 4 To save the changes made to the view, click *View > Saving > Save*.

Customizing Grid Lines

By default, the views displayed by Novell ZENworks Server Management do not contain grid lines. To display horizontal and/or vertical grid lines and to select a color for the grid lines:

- 1 Click *View > Settings > Appearance*.

The Appearance dialog box is displayed.

- 2 Select the grid line style you want to use from the *Style* drop-down list. You can choose to have:
 - ◆ No grid lines (default)
 - ◆ Horizontal grid lines only
 - ◆ Vertical grid lines only
 - ◆ Vertical and horizontal lines

- 3 If you want to select a color for the grid lines, click the *Color* button.

The Color Chooser dialog box is displayed. This dialog box includes three tab pages — *Color Swatches*, *HSB*, or *RGB* — allowing three methods of color selection.

- 4 Select the color you want to use for the grid lines using one of the three tab pages, then click *OK* to close the Color Chooser dialog box.
- 5 Click *OK* to close the Appearance dialog box.
- 6 To save the changes made to the view, click *View > Saving > Save*.

Displaying the View Title

You may find it useful to display the view name at the top of the right frame to help you keep track of where you are within the Novell ZENworks Server Management Novell ConsoleOne.

To display the view title:

- 1 Click *View > Show View Title*.

22.3.2 Modifying Columns

In a tabular view, you can change the columns in the following ways:

- ♦ “Resizing Columns” on page 856
- ♦ “Adding and Removing Columns” on page 856
- ♦ “Changing the Column Order” on page 856

Resizing Columns

To resize a column:

- 1 Move the mouse pointer to the margin between the columns you want to adjust.
- 2 When the pointer changes to a sizing arrow, drag the column to the width you want.
- 3 To save the changes made to the view, click *View > Saving > Save*.

Adding and Removing Columns

To add or remove columns from a view:

- 1 Click *View > Settings > Column Selector*.
- 2 To add a column, select the column name from the *Available Fields* list, then click *Add*.
- 3 To remove a column, select the column name from the *Show These Fields in This Order* list, then click *Remove*.
- 4 Click *OK*.
- 5 To save the changes made to the view, click *View > Saving > Save*.

Changing the Column Order

To change the order in which columns are displayed:

- 1 Click *View > Settings > Column Selector*.

- 2 Select the column you want to move from the *Show These Fields in This Order* list, then click the *Move Up* or *Move Down* button to change the location of the column.
- 3 Click *OK*.
- 4 To save the changes made to the view, click *View > Saving > Save*.

22.3.3 Filtering Views

You can display the alarms in a tabular view based on filter conditions. The filter applies only to the current management session and clears after you exit Novell ConsoleOne.

You set up a filter by selecting a criteria from four drop-down lists or entering a criteria. You can either set up simple filters that require only one line, or complex filters composed of multiple lines or groups of lines. If you set up a filter using more than one line, you must also specify the logical relationship between the line and/or group of lines.

To set up a filter:

- 1 Go to the required view.
- 2 Click *View > Settings > Filter*.
- 3 Select the column by which you want to filter alarms from the first drop-down list.
- 4 Select an operator from the second drop-down list.

The operator defines the constraint value set to the column. You can specify any of the following values for the alarm display - equal to, not equal to, greater than, less than, greater than or equal to, less than or equal to, contain, or start with the value you select in the third drop-down list. The list of available operators depends on the selected column.

- 5 Select a value from the third drop-down list.
- 6 Specify how this filter statement relates to other statements you plan to define by selecting a value from the fourth drop-down list.
 - ♦ If this is the only filter statement or if it is the last statement in a group, select *End*.
 - ♦ If you want to add a line below the current filter statement, select *New Row*. A new line is added. You must define the logical relationship between the previous line and the new line. The alarms will be displayed based on the logical condition you have specified. Select *And* to satisfy both the filter conditions. Select *Or* to satisfy any one of the filter conditions for the alarm to be displayed.
 - ♦ If you want to add one or more lines that are unrelated to the preceding lines, select *New Group*. A new line is added. An additional drop-down list separates the new line from the preceding lines. Select a value from this drop-down list to indicate the relations between the filter statements. Select *And* if you want both the filter statements to be satisfied. Select *Or* if you want only one of the filter statements in one of the groups to be satisfied. Select *End* from the fourth drop-down list when you add a new group.
- 7 Click *OK* if you have finished defining filters.

The view is updated to display only those entries that meet the filter criteria you defined.

22.3.4 Sorting Views

Using the sorting feature to modify the order in which the entries in a tabular view. You can sort the entries in the following two ways:

- ♦ “[Sorting the View Using a Single Column](#)” on page 858
- ♦ “[Sorting the View Using Multiple Columns](#)” on page 858

For instructions on sorting alarms, see “[Sorting Alarms](#)” on page 927.

Sorting the View Using a Single Column

To sort the entries displayed in the view by a single column:

- 1 Double-click the column header for the column by which you want to sort the entries.

When you double-click the column header, the entries in the view are sorted by that column in descending order (the most recent entries first). To sort the entries by ascending order (oldest entries first), double-click the column header again.

Sorting the View Using Multiple Columns

To sort the view using multiple columns:

- 1 Click *View > Settings > Sort*.
- 2 Select the first column you want the entries sorted by from the *Sort Items By* field.
- 3 Select the appropriate radio button to indicate whether you want the entries sorted in ascending or descending order.
- 4 Select the second column by which you want entries sorted from the *Then By* field, then click the *Ascending* or *Descending* radio button to specify the sort order.
- 5 Repeat [Step 4](#) for each subsequent column for which you want entries sorted.
- 6 Click *OK*.

The entries are now sorted according to the criteria you specified.

22.3.5 Printing a View

To print a view:

- 1 Go to the view you want to print.
- 2 Click *File*, then click *Print*.
- 3 In the Print dialog box, select the print options you want, then click *OK*.
- 4 In the next Print dialog box, click *OK*.

22.3.6 Exporting a View

You can export a tabular or graphical view to one of the following file formats:

- ♦ HTML
- ♦ Comma-delimited text files (.csv)

- ♦ Tab-delimited text files (.txt)
- ♦ Blank-space-delimited text files (.txt)

To export a view:

- 1 Go to the view you want to export.
- 2 Click *File*, then click *Export*.
- 3 From the *Export File Type* drop-down list, select the format to export the view.
- 4 Enter the path and name of the file you want to save in the *Filename* field or click *Browse* to search for a location you want to export the file to.
- 5 Click *OK*.

22.3.7 Saving Views

By default, any of the changes you make to the appearance, content, sorting, or filtering of a view are discarded when you exit Novell ConsoleOne. If you want to retain the changes you have to explicitly save the view.

This section includes the following topics:

- ♦ “Saving the Existing View” on page 859
- ♦ “Creating a New View” on page 859

Saving the Existing View

If you want to permanently modify the existing view to reflect the changes you made, you can simply save the view as follows:

- 1 Modify the view as desired.
- 2 Click *View > Saving > Save*.

The next time you display the view, the changes will be retained.

Creating a New View

In some cases, you might find it useful to create a new view with the changes made. The existing view is left unmodified and you can save the new view under a different name.

To save the view under a new name:

- 1 Modify the view as desired.
- 2 Click *View > Saving > Save As*.
- 3 Enter a name for the view in the *Enter New View Name* field, then click *OK*.

22.3.8 Deleting and Renaming Custom Views

To rename or delete the custom views you have saved:

- 1 Click *View*, then *Saving*.
- 2 To rename a custom view, select the view from the *Saved Views* list, then click *Rename*.

or

To delete a custom view, select the view from the *Saved Views* list, then click *Delete*.

- 3 When you have finished modifying your saved views, click *Close*.

22.3.9 Displaying Multiple Views in Novell ConsoleOne Views

The *View* in a *New Window* option enables you to display multiple tabular views, trend views, and composite views in the Novell ConsoleOne Views window.

A *View in New Window* icon , identifies the views that you can display in the Novell ConsoleOne Views.

If you have multiple views in the Novell ConsoleOne Views, refreshing will take time.

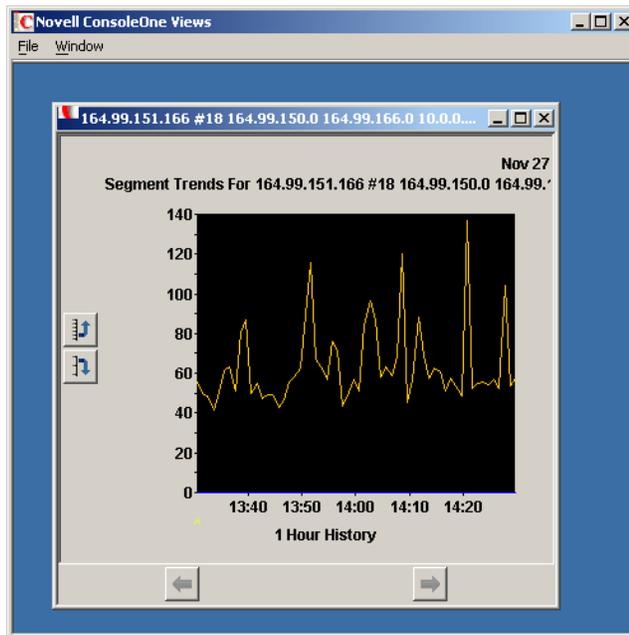
IMPORTANT: You cannot display all the views in the Novell ConsoleOne Views.

To display a view in the Novell ConsoleOne View window:

- 1 Select a view in Novell ConsoleOne.
- 2 Select *File*, then *View in New Window*.

or

Click .



You can rename the view, and also arrange the view as a tile or a cascade.

To rename a view:

- 1 In Novell ConsoleOne Views, click *File*, then *Rename*.
- 2 Specify a new name, then click *OK*.

To tile or cascade the views:

1 In Novell ConsoleOne Views, click *Window*, then *Tile*.

or

Click *Window*, then *Cascade*.

Understanding Network Discovery and Atlas Management

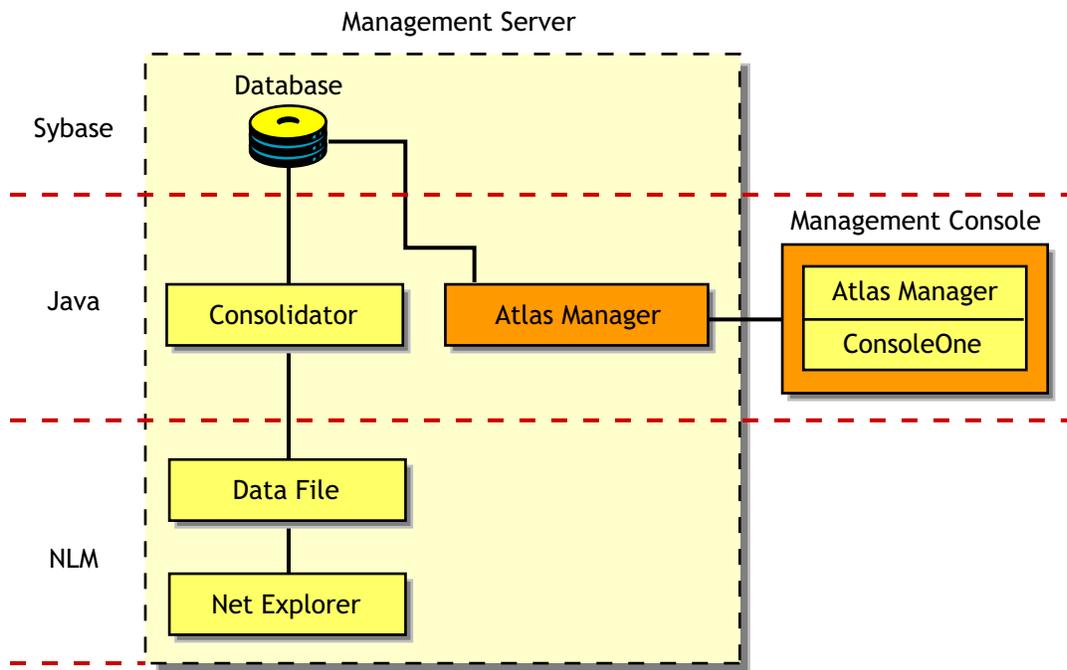
23

Discovery is the process of determining the topology of your network. You can manage, monitor and display the components of your network from Novell® ConsoleOne®. Discovery involves the following three major components of the Novell ZENworks® Server Management software:

- ♦ **Discovery software:** A set of Novell NetWare Loadable Module™ (NLM™) files that run on a management server and discovers the network topology
- ♦ **Consolidator software:** Software that runs on the management server, which reads the data discovered by discovery, and populates the Topology database.
- ♦ **Atlas Manager software:** Software that reads the Topology database, creates an atlas database, and displays the network topology in an atlas on Novell ConsoleOne.

Figure 23-1 shows a high-level view of the discovery components:

Figure 23-1 Discovery Components



This section deals with the following topics:

- ♦ Section 23.1, "Understanding Network Discovery," on page 864
- ♦ Section 23.2, "Setting Up Discovery," on page 897
- ♦ Section 23.3, "Managing the Atlas," on page 908

23.1 Understanding Network Discovery

The NetExplorer™ software drives the discovery process on the management server. The discovered information is populated in the Topology database. The Atlas Manager creates a related atlas database which encapsulates the topology information and adds information related to how the user views the maps.

The following sections will help you understand the network discovery process:

- ◆ [Section 23.1.1, “Discovery Components,” on page 864](#)
- ◆ [Section 23.1.2, “Discovery Process,” on page 874](#)
- ◆ [Section 23.1.3, “What Is Discovered,” on page 881](#)
- ◆ [Section 23.1.4, “File-Based Discovery,” on page 890](#)
- ◆ [Section 23.1.5, “Discovery Console,” on page 892](#)
- ◆ [Section 23.1.6, “Effects of Discovery on Maps,” on page 895](#)

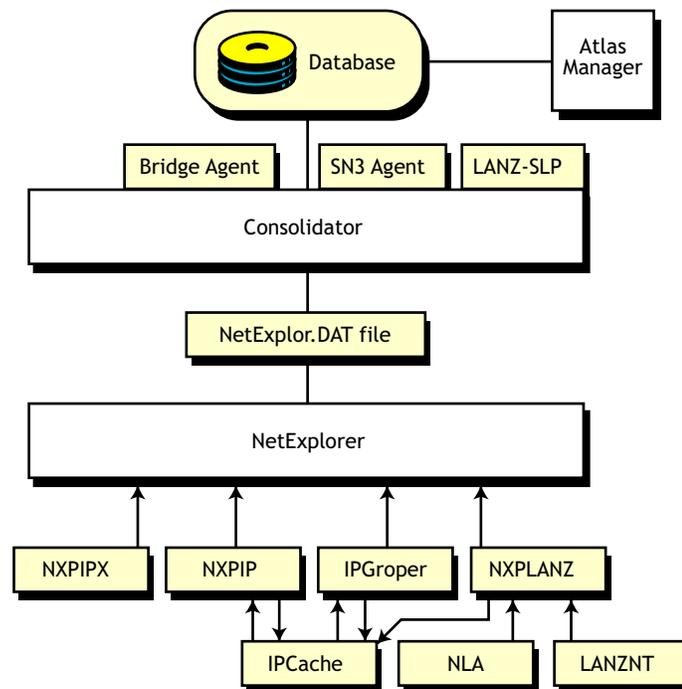
23.1.1 Discovery Components

The NetExplorer and Consolidator software that runs on the management server aids in discovering your network and updating the database.

Your network is automatically discovered by NetExplorer when you start it for the first time.

[Figure 23-2](#) shows the discovery components on the server:

Figure 23-2 *Discovery components on the Server*



The NetExplorer system consists of the following interdependent components:

- ◆ “Discovery” on page 865
- ◆ “Consolidator” on page 867
- ◆ “Atlas Manager” on page 868
- ◆ “Database Object Editor” on page 869
- ◆ “Management Console Software” on page 873
- ◆ “Additional Novell ZENworks Server Management Components” on page 873

Discovery

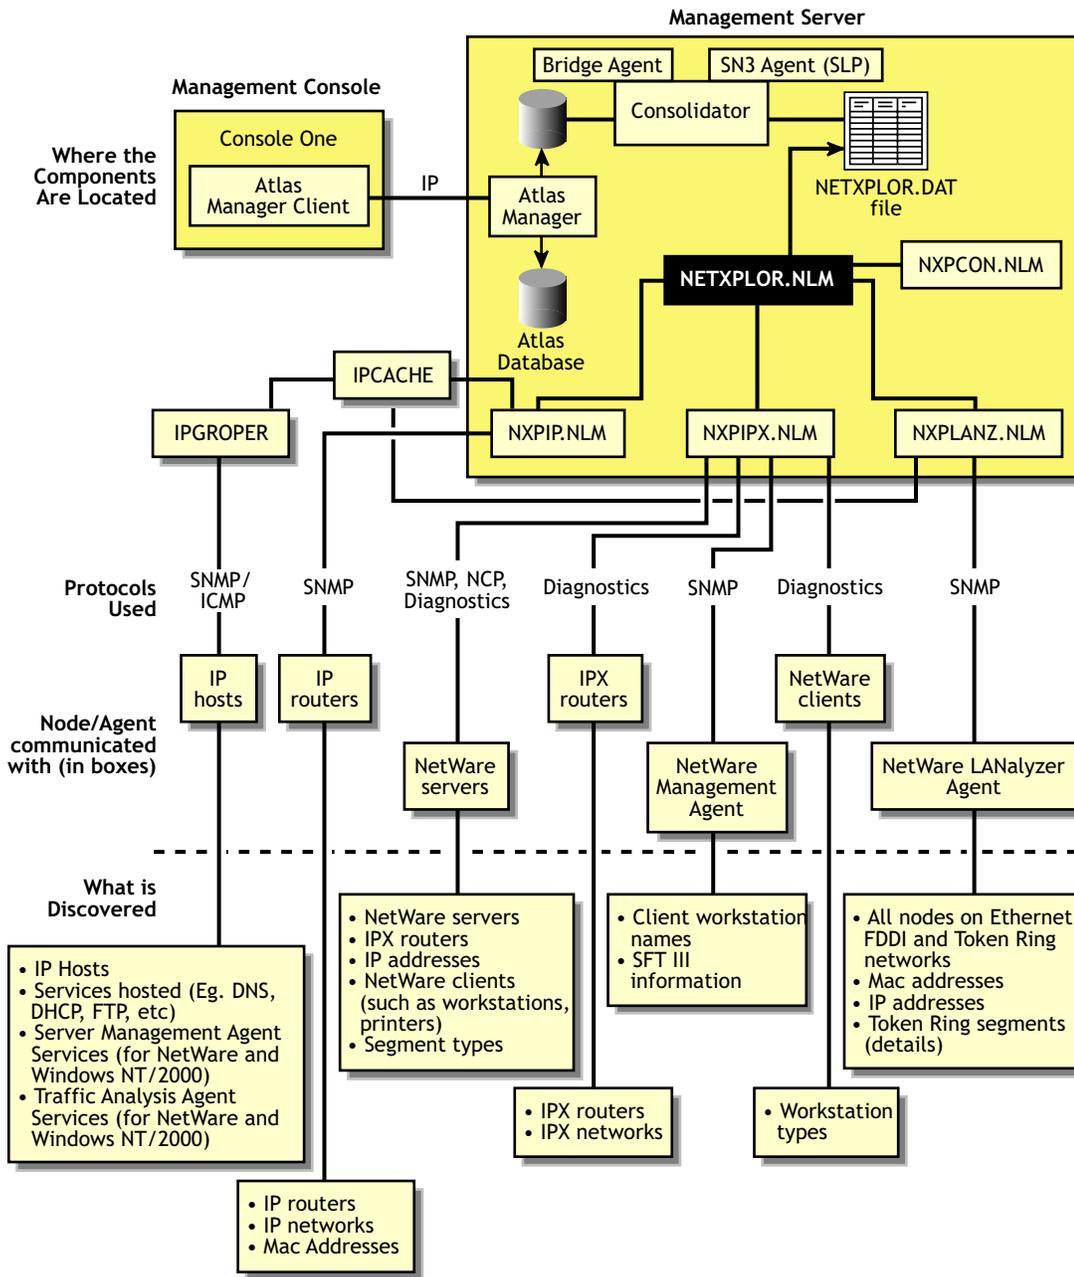
The discovery software resides on the management server and uses the discovery NLM™ software to discover the various network devices.

- ◆ `Nxpip.nlm` discovers IP routers on IP networks and sends IP router information to discovery. It communicates with the IPCACHE module to share this information with IPGROPER.
- ◆ IPGROPER detects IP host addresses and the following services: Domain Name System (DNS) names, Dynamic Host Configuration Protocol (DHCP) services, Telnet, Hypertext Transfer Protocol (HTTP), Simple Mail Transfer Protocol (SMTP), and File Transfer Protocol (FTP).
- ◆ `Nxpipx.nlm` discovers various NetWare® systems on IPX™ networks and sends information about systems to NetExplorer.
- ◆ `Nxplanz.nlm` communicates with Traffic Analysis Agents for NetWare and Windows* to gather information about all systems communicating on the segments that are monitored, and sends this information to discovery.

Figure 23-3 illustrates the architecture of the discovery system and shows the roles of the various components, network systems, and agent software.

IMPORTANT: Discovery uses the server and traffic management agents to obtain certain discovery information. Though not required, using these agents across your network enhances the accuracy and detail of logical maps displayed by Novell ConsoleOne.

Figure 23-3 Purser of the Discovery system



Supported Protocols

Novell ZENworks Server Management software supports the Service Location Protocol (SLP) on NetWare networks to enhance the discovery speed.

The server management and Traffic Analysis Agents for NetWare use the Service Advertising Protocol (SAP) to identify themselves to other components. SAP filtering prevents routers from passing SAP packets. To enable the management server and Novell ConsoleOne to receive the SAP packets that identify manageable servers, Hub Management Interface (HMI) hubs, and other servers, configure the router that is filtering SAP packets to list the specific SAP numbers that it should pass.

NetWare systems and ZENworks Server Management components use the SAP numbers listed in [Table 23-1](#).

Table 23-1 SAP Numbers used by the NetWare systems and ZENworks Server Management

| Component | SAP Number (Decimal) | SAP Number (Hexadecimal) |
|-------------------------------------|----------------------|--------------------------|
| NetExplorer NLM | 567 | 237 |
| Novell Server Management Agent | 635 | 27B |
| ManageWise Agent for Windows server | 651 | 28B |
| Traffic Analysis Agent for NetWare | 570 | 23A |
| Print server | 7 | 7 |
| Novell NetWare file server | 4 | 4 |

Consolidator

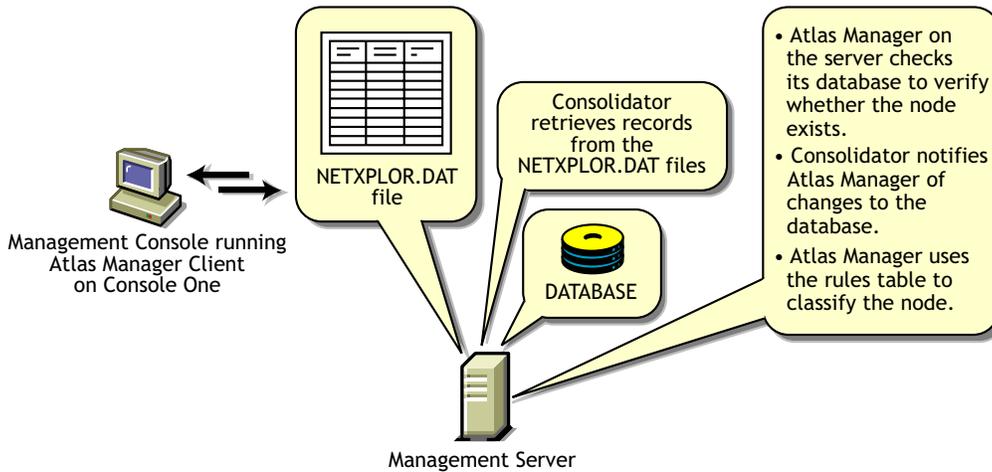
The Consolidator software resides on the management server and performs the following tasks:

- ◆ Reads the NetExplorer data files, which contains all the discovered information.
- ◆ Interprets the records in the `netexplor.dat` file.
- ◆ Checks whether the system has already been created in the Topology database. If the system does not exist in the Topology database, the Consolidator creates the system.
- ◆ Uses the Bridge agent to query the Bridge Management Information Base on IP networks and discovers which systems are connected to a port of a bridge.
- ◆ Uses the SN3 agent to get the Novell eDirectory name of NetWare servers. The SN3 agent enhances the performance of discovery by using SLP to discover NetWare servers.
- ◆ Runs the `mibcompiler.rule` file on all the discovered devices and verifies for the MIBs mentioned in the rule file on these devices and updates the database. You can also add or delete the MIBs in the `mibcompiler.rule`.
- ◆ Writes discovery information to the Novell ZENworks Server Management database.

[Figure 23-4](#) shows the tasks of the Consolidator. `Netexplor.nlm` creates the `netexplor.dat` file and the Consolidator starts reading the records from the file. If NetExplorer processes are restarted, the `netexplor.dat` file is re-created and the Consolidator requests the first record in the new file.

When the Consolidator retrieves a record from the `netexplor.dat` file, it searches for the record in the database. If the system is not in the database, the Consolidator inserts it and notifies the Atlas Manager of the update.

Figure 23-4 Consolidator Tasks



Command Line Options

If you want to manually operate the Consolidator, use the command line options shown in [Table 23-2](#).

Table 23-2 Command Line Options for the Consolidator

| Option | Allows the Consolidator to |
|----------------------------|--|
| -notify | Notify the Atlas Manager that it has updated the database. |
| -database <i>data_path</i> | The specific location of the database file to perform operations on. |

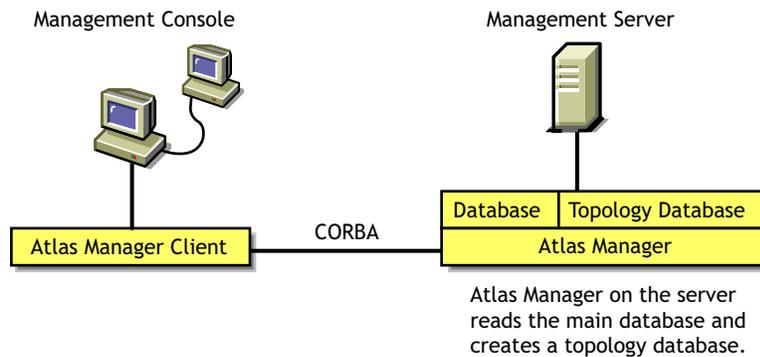
Atlas Manager

The Atlas Manager software consists of a server and a client component. The server component resides on the management server along with the Novell ZENworks Server Management topology database. The Atlas Manager server component retrieves discovery data from the topology database and creates its own atlas database.

The client component of the Atlas Manager resides on Novell ConsoleOne. The server component can communicate with several console components at any given time. Any changes made to the maps on the console (for example, rename, import, and layout) are communicated to the Atlas Manager server component, to update the atlas database.

[Figure 23-5](#) shows the Atlas Manager server and client software:

Figure 23-5 *Atlas Manager server and client software*



The Atlas Manager looks for a rule in its rules table to help classify the system. The rules help the Atlas Manager make decisions, such as which icon should be used to display the system on the maps. If the system in the record matches one of the rules, the Atlas Manager updates the database according to the rule.

Database Object Editor

The Database Object Editor supplements the discovery system. Sometimes auto discovery might not discover devices on your network, or might display incorrect information of the devices on your network. You can use Database Object Editor to add the missing entities into the database or to edit incorrect information of the entities.

The Database Object Editor client uses ConsoleOne snap-in to display the user interface. Using the Database Object Editor, you can perform operations on a segment or a node.

The Database Object Editor Server interacts with the Consolidator to process information related to the node and segment object and populates the topology database with this information.

You can use the Database Object Editor to add or delete a segment or a node and modify the segment or the node information.

To add a segment or a node:

- 1 In ConsoleOne, select *Atlas*, then click *Tools > Database Object Editor > New*.
- 2 Enter the details for the segment or the node.
- 3 Click *OK*.

To edit the information about the segment or the node:

- 1 In ConsoleOne, select the segment or the node you want to edit.
- 2 Click *Tools > Database Object Editor > Edit*.
Modify the required information.
- 3 Click *OK*.

To delete the segment or the node:

- 1 In ConsoleOne, select the segment or the node you want to delete.
- 2 Click *Tools > Database Object Editor > Delete*.

To add or delete list of MIBs to a node:

- 1 In ConsoleOne, select the node where you want to add/delete the MIBs.
- 2 Click *Tools > Database Object Editor > Edit*.
- 3 Select Implemented MIBs on the left pane to add or delete MIBs.

To add or delete list of services to a node:

- 1 In ConsoleOne, select the node where you want to add/delete services.
- 2 Click *Tools > Database Object Editor > Edit*.
- 3 Select *Services* on the left pane to add or delete services.

NOTE: If you install the LANalyzer service on a node, the RMON service is also automatically installed because the LANalyzer service cannot function without the RMON service. The Service pane of the Database Object Editor displays LANalyzer service only.

To add, delete, or modify the interfaces of a node:

- 1 In ConsoleOne, select the node whose interfaces you want to add/delete/modify.
- 2 Click *Tools > Database Object Editor > Edit*.
- 3 Select *Interface Summary* on the left pane to add/delete/modify interfaces.

To modify operating system of a node:

- 1 In ConsoleOne, select the node whose operating system you want to modify.
- 2 Click *Tools > Database Object Editor > Edit*.
- 3 Change the type of Operating System in the *Operating System* field.

To add, delete, or modify the switch port to the end node connectivity:

- 1 In ConsoleOne, select the switch whose connectivity needs to be added/deleted/modified.
- 2 Click *Tools > Database Object Editor > Edit*.
- 3 Select *Switch Summary* on the left pane to add/delete/modify switch port to node connectivity.

Working with Unnumbered Links

Unnumbered links are point to point links between routers with no IP address bound to interfaces on the two ends of the WAN link.

The Router Discovery Module (NXPIP) discovers unnumbered WAN links between routers using the following methods:

- ◆ **Auto Discovery of Unnumbered Links:** In this method, the router tables for each router is obtained. The unnumbered links between two routers are identified by correlating the routers that point to each other.
- ◆ **Manual Configuration of Unnumbered Links:** You can use this method of discovery in scenarios where auto discovery fails to discover a link or incorrectly discovers a link. Using the `nxpip.ini` file you can specify a link between two routers having unnumbered interfaces or prevent auto discovery from creating a link. Refer to the `nxp.ini` file in the `installation_directory\novell zenworks\mms\mwserver\nmdisk` directory for configuration details.

If the interface type of the unnumbered link that is discovered is not a PPP, ATM, FrameRelay, or X 25, the unnumbered link will be created with the interface type as unknown.

To enable the unnumbered link discovery, make the following changes in the `netexplor.ncf` file:

- 1 Open the `netexplor.ncf` file in the `installation_directory\novell zenworks\mms\mwserver\nmdisk directory`.
- 2 Locate the line `Load NXPIP`. If you are unable to find this line, use `NXPCON` to enable the `NXPIP` module. This line will now be added to the `netexplor.ncf` file.
- 3 Add the following options: `/autould/iniuld`. The line should now read as follows:
`Load NXPIP /autould/iniuld`.

NOTE: Add any one of the numbered IP address of the routers to the additional routers list. For more information on how to add the additional routers, see [“Specifying a Seed Router and Additional IP Routers” on page 906](#).

When you restart NetExplorer, the unnumbered links will be discovered and displayed in the atlas namespace.

Viewing the Unnumbered links in the Atlas Namespace

To view the unnumbered links on a router:

- 1 Go to the properties page of a router.
- 2 Select the *Computer Attributes* tab.
- 3 The *IP Address* field displays the string Unnumbered IP Address (*n*), where *n* is the number of unnumbered interfaces of the router.

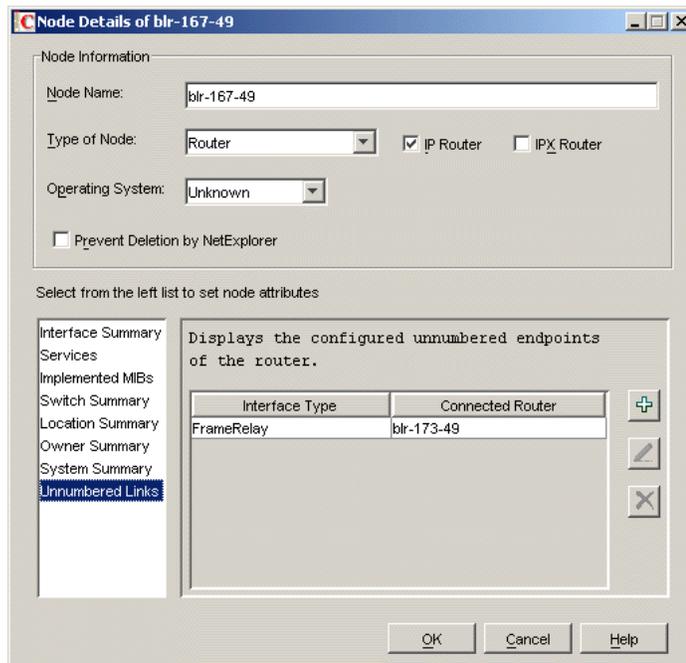
To view the unnumbered links on a segment:

- 1 Go to the properties page of a segment.
- 2 Select the *Segment Attributes* tab.

The Unnumbered Network string is used in the *IP Address* field to display the network number of the unnumbered link.

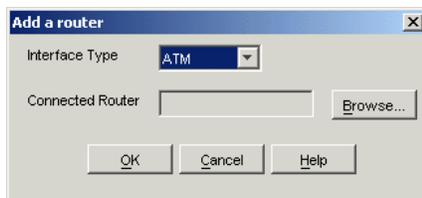
Using the *Unnumbered Links* tab in the Database Object Editor you can add, modify, or delete the unnumbered links for the router. The unnumbered links includes the interface type and the connected router. All the unnumbered links you created and configured will be displayed in the list.

Figure 23-6 *Unnumbered Links*



Adding an Unnumbered Link

- 1 In the *Unnumbered Links* tab, click *Add*.



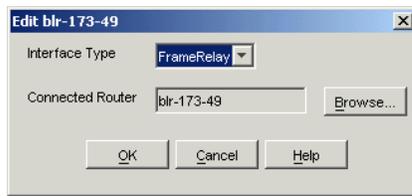
- 2 In the *Interface Type* drop-down list, select the interface type of the router.
- 3 Click *Browse* to select the router that you want to connect.



- 4 Click *OK*.

Editing an Unnumbered Link

- 1 In the *Unnumbered Links* tab, select the unnumbered link you want to edit from the list, then click *Edit*.



- 2 In the *Interface Type* drop-down list, modify the interface type of the router.
- 3 Click *Browse* to select the router that you want to connect.
- 4 Click *OK*.

Deleting an Unnumbered Link

- 1 In the *Unnumbered Links* tab, select the unnumbered link you want to delete from the list.
- 2 Click *Delete*.

If you delete any of the connected node, the corresponding unnumbered link will be deleted. All the information pertaining to the unnumbered link will be updated in the database.

Management Console Software

The management console software snaps in to Novell ConsoleOne. Management sites are created in Novell ConsoleOne. In each site, an atlas is created that maintains the integrity of the discovery information.

Additional Novell ZENworks Server Management Components

The `nxpip.nlm`, `nxpipx.nlm`, and `nxplanz.nlm` software operate in conjunction with the following components:

- ♦ “Traffic Analysis Agent for NetWare Servers” on page 873
- ♦ “Server Management Agent for NetWare Servers” on page 874
- ♦ “Bindery of Novell NetWare Servers” on page 874

Traffic Analysis Agent for NetWare Servers

The Traffic Analysis Agent for NetWare is a set of NLM files that provides traffic analysis of Ethernet, Fiber Distributed Data Interface (FDDI), or token ring segments. The Traffic Analysis Agent discovers all systems on the segments it monitors, regardless of the protocols the systems use. You can monitor multiple segments by placing agents on each segment.

The `nxplanz.nlm` software on the management server uses SNMP to query servers running the Traffic Analysis Agent for information about each system that resides on their segments.

IMPORTANT: For an effective discovery process, you should have the Traffic Analysis Agent monitoring each source-routed token ring segment.

Server Management Agent for NetWare Servers

To discover IPX servers and workstations, managed servers are any NetWare servers with the server management agent installed. Server management agents respond to SNMP queries from `nxpipx.nlm` with the username and address of those workstations that are logged in to the server. `Nxpipx.nlm` obtains SFT III™ server information from the server management agent. For effective results, you should install a management agent on every NetWare server on your network.

Bindery of Novell NetWare Servers

`Nxpipx.nlm` queries all NetWare servers for information in their binderies. All NetWare servers allow their binderies to be examined by the discovery process when their security settings are set to the default values.

For the NetExplorer NLM software to discover the login names of workstations attached to a NetWare server, a server management agent must be installed on the server.

23.1.2 Discovery Process

NetExplorer discovers your network continually. The following sections discuss the discovery processes:

- ♦ [“Discovery Cycles” on page 874](#)
- ♦ [“Continuous Discovery” on page 878](#)

Discovery Cycles

When you first start discovery, you should let it run as long as necessary to build the baseline data. Very small networks might take one or two hours, while very large networks (several thousand nodes) might require a day or two to be discovered.

The discovery process occurs in cycles. A cycle is the process by which a discovery module identifies every node it can at a time. You can configure discovery on the server to discover only certain addresses, thus reducing the duration of a cycle. For more information, see [“Changing the Discovery Scope” on page 903](#).

The initial cycle continues until no additional devices are discovered. This initial cycle gathers information that might be insufficient to classify certain devices or to identify the correct segment for each device. Further discovery cycles provide additional, new, and changed information. As discovery cycles proceed, the information becomes more accurate.

Each discovery process queries the network using different methods to discover systems. Four independent discovery modules run in the order mentioned below during each discovery cycle:

1. IP router discovery on IP networks only.

This process, run by the NXPIP module, starts from the local router. Using the local router's routing table information, NXPIP discovers other routers on the network. It then uses the routing table information to further discover the network. This process is repeated for each router discovered.

The NXPIP module stores the router address information and information about any IP-bound network device in the IPCACHE module.

`Nxpip.nlm` is installed on the management server. It uses SNMP to discover IP routers. To use this NLM, your management server must also be running TCP/IP bound to at least one of your network's interface boards. `Nxpip.nlm` uses MIB-II information, such as the system table, routing table, interface table, interface data-link type and frame type, and segment data-link type. Note that because there are different versions of MIB-II implementations for different vendors, the information you receive might differ.

IMPORTANT: If you have specified an additional level of control by allowing certain IP addresses to perform SNMP queries to the routers, ensure that the IP address given to the Novell ZENworks Server Management server is privileged to query all the routers in the network. Otherwise, discovery will not be complete, and incomplete network information will appear in the Islands page of the atlas.

2. IP discovery of workstations and servers.

This process, run by the IPGROPER module, receives the router and network information written into the IPCACHE by the NXPIP module as the input. RMON, based discovery run by the NXPLANZ module also writes the information about the networks and IP hosts that it discovers into IPCACHE. This also acts as an input to the IPGROPER module.

It queries each router that has been discovered by NXPIP for its ARP tables, identifying each active IP host on the network. For IP addresses that are not found in the ARP table of any of the routers, IPGROPER tries to ping and identify whether a host by that IP address is alive.

IPGROPER queries each IP host that is identified to be alive for information about the following hosted services: HTTP, DHCP, Telnet, SMTP, and DNS. It also verifies whether the server management software and the Traffic Analysis Agents are installed and running on this host.

Simultaneously, the IPGroper module queries the DNS server specified in the `sys:\ect\resolv.cfg` file on the management server for the DNS names of all these IP hosts.

IMPORTANT: For a server or a segment to be manageable, it is important to discover the server management agent and the Traffic Analysis Agent running on an IP host on that server or the segment.

3. IPX discovery on all networks, including NetWare/IP networks:

This process, run by the NXPIPX module, starts at the management server itself to discover its IPX address, the LAN type of each adapter, and SAP information about other known devices and their services. After gathering this information, NXPIPX requests the same types of information from each device listed in the bindery. This process is repeated each time NXPIPX discovers a new device.

`Nxpipx.nlm` uses a variety of NetWare, SNMP, and IPX protocols, such as IPX diagnostics, to discover NetWare servers, IPX routers, and IPX workstations.

IMPORTANT: When `nxpipx.nlm` is loaded, a working directory named NXPWORK is created by default under the `installation_volume\install_dir\Novell ZENworks\mms\mwserver\nmdisk` subdirectory. During installation, you can specify a different path to create the NXPWORK subdirectory. NXPIPX puts all of its temporary files in this directory. Do not read, modify, or delete any file in this directory because this might cause some discovery process to not function.

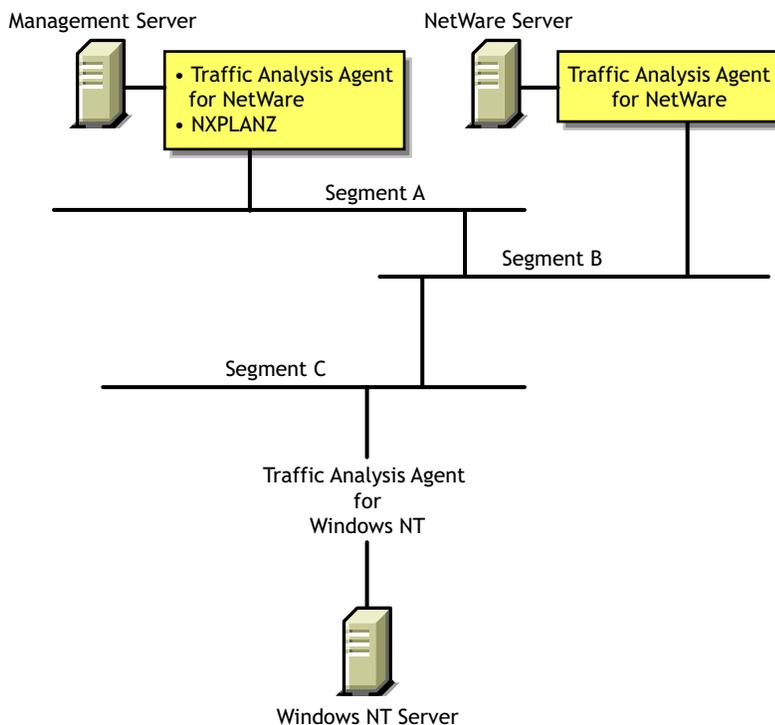
4. RMON based discovery of IP Hosts.

This process, run by the NXPLANZ module, starts by identifying all the remote agents, which includes the Traffic Analysis Agents for NetWare and Windows. The Traffic Analysis Agents on a segment discover devices based on the IP address to MAC address binding data contained in packets that are transmitted on the segment. The NXPLANZ module on the management server retrieves the data by using SNMP to communicate with the Traffic Analysis Agents.

The NXPLANZ module reports information about the Traffic Analysis Agents on your network and the IP hosts on the segments monitored by these Traffic Analysis agents to NetExplorer. The information about the networks monitored by the Traffic Analysis Agents and IP hosts on the monitored networks is also written to IPCACHE to enhance the effectiveness of service discovery by the IPGROPER module.

Figure 23-7 shows NXPLANZ querying Traffic Analysis Agents software on segments B and C, respectively.

Figure 23-7 NXPLANZ Querying Traffic Analysis Agent Software



To improve the effectiveness of the discovery, ensure that the Traffic Analysis Agent is installed and running on each network segment that you want to discover. If SLP is disabled on your network or if SAP packets are filtered by the routers in your network, NXPLANZ may not be able to discover all the Traffic Analysis agents in the network.

In order to ensure that all the Traffic Analysis agents on your network are being queried by the NXPLANZ or NXPLANZ module, specify these Traffic Analysis agents explicitly using NXPCON.

During the initial discovery cycle, these modules run sequentially. As a result, information about the Traffic Analysis Agent software is discovered late.

In later discovery cycles, the four modules run concurrently. They continue their discovery processes, but send only new or changed data to `netexplor.nlm`. As additional data arrives, segments can be consolidated, devices can be placed on the appropriate segments, and new devices can be discovered.

Each succeeding cycle of different discovery NLM files has the potential to provide key information that finally identifies a device and provides sufficient data for NetExplorer to consolidate the data.

The data discovered by the NLM processes is communicated to Novell ConsoleOne through the Atlas Manager. **Figure 23-8** shows the relationship of the discovery NLM processes, NetExplorer, and Novell ConsoleOne. See **“Discovery Process”** on page 874 for a description of how these pieces operate together to discover the contents and topology of a network.

Figure 23-8 The relationship of the discovery NLM processes, NetExplorer, and the management console

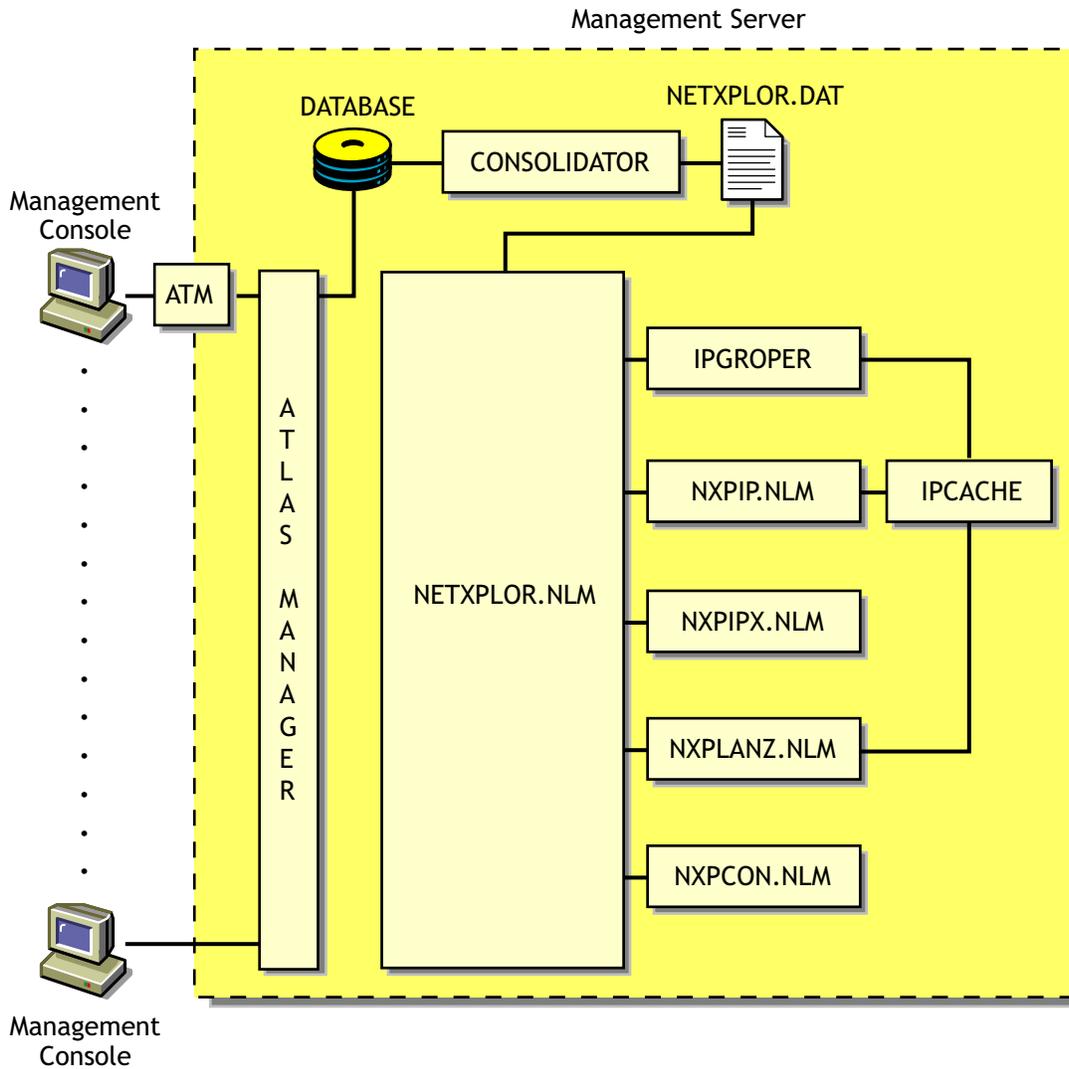


Table 23-3 summarizes the default seed and scope and user-definable changes for each discovery module:

Table 23-3 Summary of the default seed, scope and user-definable changes for each discovery module

| Discovery Module | Default Seed Information | Default Scope | User-Definable Changes |
|------------------|--|--|--|
| NXPIP | Examines the management server routing table. Places the router addresses in the IPCACHE module. | Entire network if community string matches. | Reduce scope by specifying IP scope information in NXPCON. If public SNMP community string is not used, list SNMP community strings of routers in NXPCON. |
| IPCACHE | Supporting module in NetExplorer. Contains temporary information about devices and networks which is used by NXPIP, IPGROPER and NXPLANZ. | | |
| IPGROPER | <ol style="list-style-type: none"> 1. Queries each router address in IPCACHE for ARP tables to identify network devices. 2. Queries each network device for the services it hosts (FTP, HTTP, Telnet, SMTP, DNS, and DHCP) and their DNS names. 3. Discovers hosts running server management and Traffic Analysis Agents. | All IP networks connected to routers already discovered by NXPIP | <ul style="list-style-type: none"> ◆ Enable or disable autodiscovery ◆ Enable or disable file-based discovery |
| NXPIPX | Examines the management server's configuration. | Entire IPX internetwork. | Reduce scope by specifying IPX scope information in NXPCON. |
| NXPLANZ | Examines the list of servers running Traffic Analysis Agent software listed in NXPCON. | All segments with Traffic Analysis Agent software. | Specify name and IP addresses of Traffic Analysis Agent for Windows in NXPCON. If SLP is disabled or SAP is being filtered, specify the name and address in NXPCON for the Traffic Analysis Agent for NetWare. |

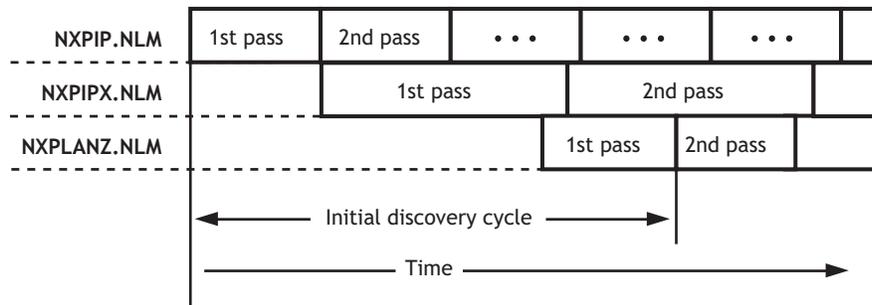
Continuous Discovery

NetExplorer discovers the internetwork on which it resides, through a process initiated and controlled by `netexplor.nlm`. Initially, each discovery NLM identifies itself to `netexplor.nlm`, which then begins the initial discovery cycle. The cycle starts with NXPIP discovery, followed by NXPIPX discovery, and finally NXPLANZ discovery. The discovery cycles of IPGROPER are not controlled by `netexplor.nlm`. After starting, it runs continuously. Information gathered by NetExplorer is stored in the `netexplor.dat` file on the management server.

In [Figure 23-9](#), each of the discovery processes is shown in relationship to time. After NXPIP finishes its first pass, NXPIPX begins and NXPIP starts over. After NXPIPX finishes its first pass, NXPLANZ begins and NXPIPX starts its second pass.

Unless otherwise directed, all three of the discovery processes run continually to detect changes to the network. Any changes to the network are saved as records in the `netexplor.dat` file. When all three discovery processes have completed one pass, the initial discovery cycle is complete.

Figure 23-9 *The discovery processes in relationship to time*



The following sections describe each sequence in greater detail:

- ◆ [“NXPIP” on page 879](#)
- ◆ [“NXPIPX” on page 879](#)
- ◆ [“NXPLANZ” on page 880](#)
- ◆ [“IPGROPER” on page 880](#)
- ◆ [“NETXPLOER” on page 880](#)
- ◆ [“SNMP Community String Discovery” on page 880](#)

NXPIP

The first sequence in the NetExplorer discovery cycle involves the discovery of IP routers. NXPIP locates its local router using TCP/IP configuration information. NXPIP then queries the router for the identity of other routers on the network. NXPIP queries the MIBs on the routers using SNMP to collect the IP addresses, interface types, and MAC addresses.

By default, NXPIP attempts to discover your entire IP network. You can restrict the scope of the IP discovery by specifying the scope information in NXPCON.

NXPIPX

NXPIPX uses a series of techniques, including SNMP, RIP, IPX, and SPX™ diagnostics to discover the attached IPX or NetWare/IP internetwork. After NXPIP completes its first pass, NXPIPX begins discovery at the management server. NXPIPX examines its own server and discovers the names of other servers. It then queries each of these servers to discover more servers and repeats this process until no more servers are found.

In addition, NXPIPX reads the connection table of each NetWare server to determine which NetWare clients are logged in to the server. NXPIPX sends IPX diagnostic packets to each client to collect additional information. NXPIPX will not discover clients that do not appear in the connection table because they have not been logged in recently and clients whose diagnostics are turned off. It is therefore important to leave IPX diagnostics enabled on NetWare clients.

NXPIPX also discovers IPX routers in your network. Third-party IPX routers are discovered only if there is a NetWare server on the routed segment. NXPIPX does not discover interface information when routed segments do not have NetWare servers.

By default, NXPIPX attempts to discover your entire IPX internetwork. You can restrict the scope of discovery by specifying a list of IPX network numbers using NXPCON. For NXPIPX to discover other IPX nodes ensure that one of the IPX numbers is bound to the management server.

NXPLANZ

The Traffic Analysis Agent for NetWare monitors every packet on the network segment it is installed on. It creates a list of physical (MAC) addresses and IP addresses of all the systems communicating on the segment on the local memory. After NXPIPX completes its first pass, NXPLANZ uses SNMP to query all servers with Traffic Analysis Agents installed to read the list of workstations communicating on the network. NXPLANZ also obtains a list of the agents running on the servers from NXPIPX.

IPGROPER

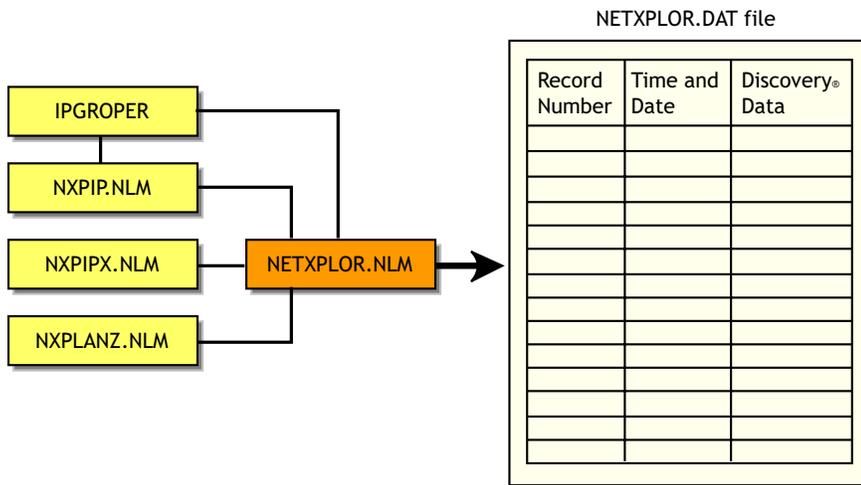
The information about the routers and network segments written into IPCACHE by NXPIP, and the information about network segments, and hosts written to IPCACHE by NZPLANZ forms the input to the IPGROPER module. For each network segment, IPGROPER tries to discover all the hosts on that network, the DNS names of the hosts, the services hosted on them, server management agents, and Traffic Analysis Agents.

NETXPLOER

As the discovery processes gather information about systems on the network, they forward packets of related data to `netexplor.nlm`. `Netexplor.nlm` places these packets, along with a record number and a time stamp, into the `netexplor.dat` file, as shown in [Figure 23-10](#).

NOTE: Discovery re-creates the `netexplor.dat` file each time you load `netexplor.nlm`. Therefore, the discovery data stored at the management server from previous runs of the NetExplorer NLM processes is not retained when you restart `netexplor.nlm`.

Figure 23-10 Discovery process of NETXPLOER.NLM



SNMP Community String Discovery

Each time NetExplorer tries to access a system through SNMP, it uses the community strings that have been configured using the NXPCON utility on the management server. When it encounters a

new system, it tries each of the configured community strings. After it has found a community string for a particular IP or IPX address, it records this name in a file so that in subsequent cycles it does not need to retry with the other configured names.

You can view these community strings using NXPCON. The community strings are used in the order specified. Therefore, the most-used community string should be configured first in the list.

IMPORTANT: An SNMP query with an invalid SNMP community string results in no response from the target system and the request times out.

23.1.3 What Is Discovered

NXPIP, NXPIPX, and NXPLANZ use a variety of techniques to discover the following categories of network objects and present them in the atlas:

- ♦ “Systems” on page 881
- ♦ “Network Segments” on page 888

Generally, information gathered by NXPIP and NXPIPX is sufficient to place systems on the network maps correctly. When NXPIP and NXPIPX have not discovered systems, NXPLANZ retrieves MAC addresses collected by the Traffic Analysis Agent software and the new systems are added to the database. Consequently, all systems are discovered on segments monitored by the Traffic Analysis Agents.

Systems

Table 23-4 shows the different types of systems discovered:

Table 23-4 *Types of systems discovered*

| System | Comment |
|------------------------------------|--|
| Novell Server Management Agent | Service type of 563 decimal (Novell Server Management Agent 1.5 or 1.6) or 635 decimal (Novell Server Management Agent 2.6) or Novell Server Management Agent MIB implemented. |
| Management Agent for Windows | Management Agent MIB implemented. |
| Traffic Analysis Agent for NetWare | Service type of 570 decimal or Traffic Analysis MIB implemented. |
| Traffic Analysis Agent for Windows | Traffic Analysis MIB implemented |
| Novell NetWare File Server | Service type of 4 (file server). NXPIPX discovers all NetWare servers. |
| Novell NetWare Print Server™ | Service type of 71 or 7 decimal. |
| IPX Router | System with more than one adapter connected to different IPX networks. |
| IP Router | System that is configured as an IP router in MIB-II (IP forwarding enabled). |

| System | Comment |
|--------------------------------------|---|
| Novell NetWare Client Workstation | System that responds to IPX diagnostics requests as an IPX workstation (has the Novell NetWare Shell loaded). |
| SFT III IOEngine | Discovered by the IPX discovery module; responds with diagnostic information. |
| SFT III MSEngine | Discovered by the IPX discovery module. |
| Network Printers | Discovered if the printer generates a well-known service type. |
| Novell NetWare Connect™ | Service type of 590 decimal. |
| Novell NetWare Communications Server | Used by the Novell NetWare for SAA* services manager products; has a service type of 304 decimal. |
| Management Server | Running discovery NLM files; has a service type of 567 decimal. |
| Any System | Any system is discovered if it is connected to a LAN segment being monitored by a Traffic Analysis Agent. |

The different types of services discovered are Telnet, HTTP, DNS, SMTP, DHCP, Routers, Novell eDirectory, SFTIII, and SNMP.

The following sections contain more information about the various systems that are discovered:

- ◆ “Novell NetWare Client Workstations” on page 882
- ◆ “IP Routers” on page 883
- ◆ “Novell NetWare SFT III” on page 884
- ◆ “Systems Not Equipped with the IPX Diagnostic Responder” on page 885
- ◆ “Routers that Use Duplicate MAC Addresses” on page 885
- ◆ “Third-Party Routers” on page 886
- ◆ “Novell NetWare MultiProtocol Router with WAN Ports” on page 886
- ◆ “IPX Networks” on page 886
- ◆ “IP Networks” on page 887
- ◆ “On-Demand Links” on page 887
- ◆ “Third-Party Routers with WAN Ports” on page 887
- ◆ “Novell NetWare Connect Servers” on page 887
- ◆ “Virtual Switches” on page 887

Novell NetWare Client Workstations

NXPIPX discovers all Novell NetWare client software attached to discovered NetWare servers. Clients that are turned off or are not attached to a server are not discovered. For this reason, a NetExplorer process that is run at night or on a weekend might not yield a complete map. Note that NetWare clients must have IPX diagnostics enabled.

When you configure a NetWare client to perform a bindery login, consider the scenarios in [Table 23-5](#):

Table 23-5 Bindery Login Scenarios

| Server | Bindery Login—What Is Discovered |
|---|--|
| Novell NetWare with server management agent installed | Workstation discovered; name is discovered only if logged in with IPX as the transport for Novell NetWare 4.x and Novell NetWare 5.x |
| Novell NetWare | Workstation discovered; name is not discovered |

When you configure the client to perform a directory login, NetExplorer discovers only those systems that are logged in to an Novell eDirectory tree and not those that are merely attached to the Novell eDirectory tree. NXPIP uses SNMP community string to communicate with the management agent and query on all NetWare servers for the username.

After NetExplorer discovers a NetWare client, NXPIP queries the client using the IPX diagnostic protocol to confirm the discovery and gather more information about it. If IPX diagnostics are turned off, NXPIP does not report the system. This applies to printers as well.

IP Routers

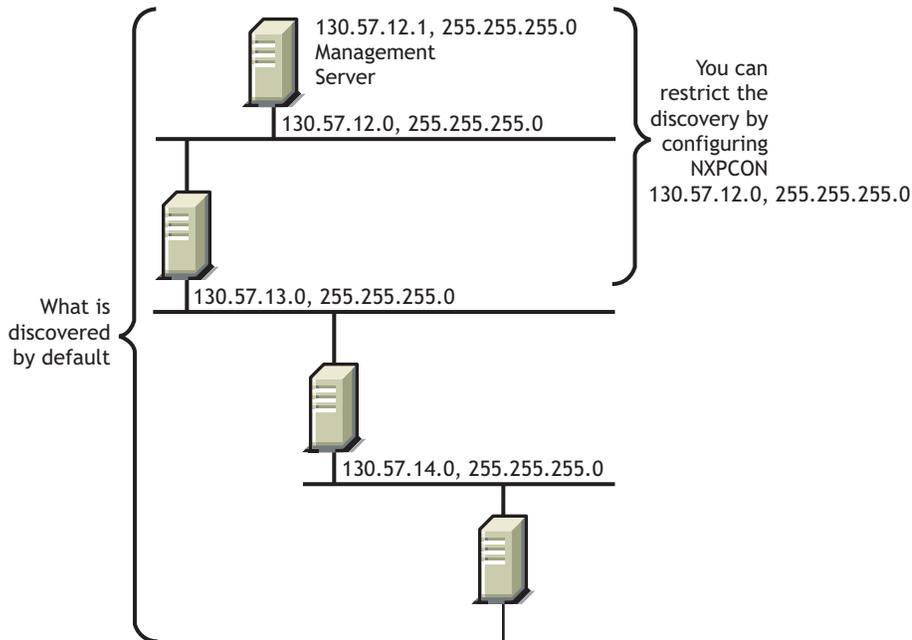
NXPIP uses SNMP to query all IP routers on the network by using the SNMP community string used by the routers. You must enter the list of community strings used by your routers using NXPCON.

You can configure this information into the router's MIB by using any SNMP configuration tool, including the SNMP MIB browser. If you configure router information such as the system name in the routers SNMP MIB, the discovery process records it in the database, allowing IP routers to be displayed with meaningful names.

By default, IP discovery discovers the entire network. The exploration can be restricted by specifying network numbers using the NXPCON Discovery Scope, and then IP Discovery Scope option. Also, if there are redundant IP routers, use the NXPCON IP Discovery, and then the IP Routers option to specify the redundant IP router address; otherwise, NXPIP does not discover it. As shown in [Figure 23-11](#), if the management server IP address is 130.57.12.0, the IP discovery NLM discovers the entire 10.57.85.0 network and its subnets.

[Figure 23-11](#) shows how Novell ZENworks for discovers IP routers:

Figure 23-11 Novell ZENworks for discovery of IP routers



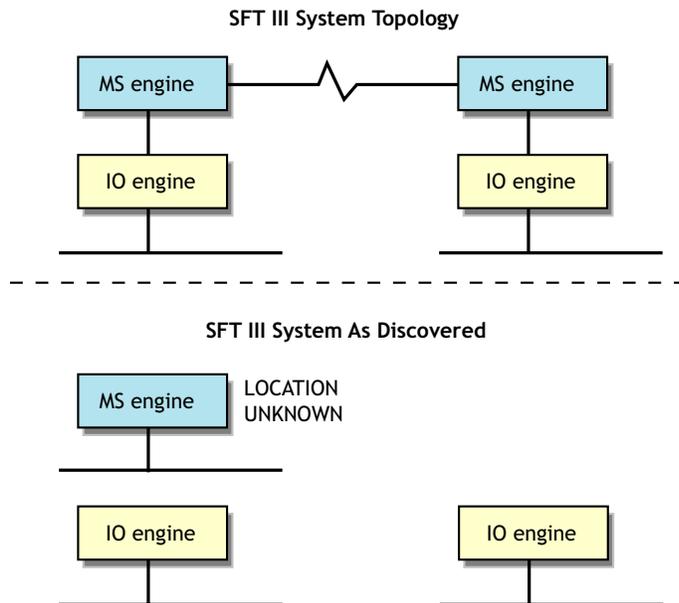
Novell NetWare SFT III

A NetWare SFT III™ server usually consists of two computer systems, each containing an input/output engine (IO engine) and a mirrored server engine (MS engine). Therefore, physically there are two IO engines and two MS engines; logically there are two IO engines and one MS engine.

If Novell Server Management Agent is loaded on an SFT III server, the MS engine and both IO engines are discovered correctly with their names and placed in the correct segment in the atlas. However, the MS engine is placed in the Islands page. This happens because the two MS engines are associated with only one logical server on the network, and the location of the MS engine might change depending on which copy of the MS engine is the primary at any given time.

Figure 23-12 illustrates NetWare SFT III server discovery:

Figure 23-12 Server discovery for NetWare SFT III



If the server management agent is not loaded on the MS engine, Discovery discovers only the MS engine and the IO engine that are primary at the time of discovery. The primary IO engine is labeled Noname in the area page. To change the name of an IO engine on a segment map, right-click the icon and click Rename.

Systems Not Equipped with the IPX Diagnostic Responder

NXPIPX discovers the following systems, but does not necessarily place them correctly in the atlas:

- ◆ NetWare for UNIX* servers
- ◆ Portable NetWare
- ◆ Access servers
- ◆ Modem servers
- ◆ Print servers

Because these systems do not respond to IPX diagnostics, they cannot answer queries from NXPIPX. Consequently, the LAN information required to place them on the maps might not be available. In this situation, NetExplorer places these systems in the Islands page of the atlas. In most cases, the presence of a Traffic Analysis Agent on each segment on which these systems appear, enables NetExplorer to obtain the missing information and correctly locate the systems in the maps.

If these systems are running IP, they will be discovered and placed correctly in the maps.

Routers that Use Duplicate MAC Addresses

NetExplorer can experience difficulties in discovering some routers because of the method routers use to identify their adapters. In some cases, the same MAC address is used on several network interfaces of a router. In these cases, it appears to NetExplorer that one adapter is connected to multiple segments. Unless otherwise specified, NetExplorer interprets multiple adapters as one adapter.

The multiple segments connected to the adapters are seen as one segment and NetExplorer consolidates the multiple segments.

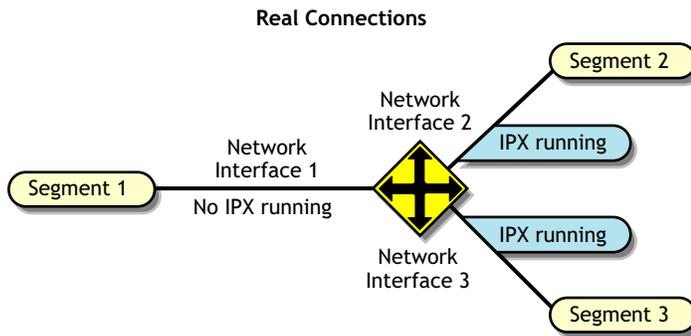
Third-Party Routers

NXPIP discovers IP-bound interfaces only. When IP is not running on a router, NetExplorer discovers the IPX-bound interfaces, which results in:

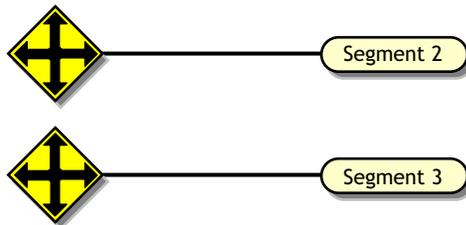
- A separate router icon is shown for each interface in the router.
- Discovered interfaces are not placed in the same router in the atlas. Therefore interconnections are incorrect on the internetwork map and the router appears as separate, multiple routers, each containing one network interface from the real router.

Figure 23-13 illustrates a router with IPX running on Network Interfaces 2 and 3 but not on Network Interface 1. NetExplorer places this router on the internetwork map as two separate systems. As shown, the connection to Segment 1 is not displayed, and the connections to Segments 2 and 3 are shown attached to two separate systems.

Figure 23-13 Internetwork map connections for an IPX router



ZENworks for Servers Internetwork Map Connections



Novell NetWare MultiProtocol Router with WAN Ports

Novell NetWare MultiProtocol Router™ (MPR) 3.0 is now bundled with Novell NetWare 5.x.

IPX Networks

Novell NetWare MPR 3.0 reports the correct segment type of the WAN links. NetExplorer detects these correctly and displays them with the appropriate icon.

IPXWAN links between Novell NetWare MPR 3.0 systems do not have an IPX network associated with them. When NetExplorer discovers such a link, it creates a name for the WAN segment of the form #UNNUM -*n*, where *n* is an integer assigned to make the segment name unique. On multi-access networks, such as frame relay and X.25, each connection in the network adds another #UNNUM -*n* to the segment name.

IP Networks

With Novell NetWare MPR 3.0, you can configure both numbered and unnumbered IP links. NetExplorer discovers numbered links correctly. NetExplorer does not discover unnumbered IP links, resulting in the Islands page.

If IP is running on a third-party router and NXPIP is running on the management server, NetExplorer discovers only the IP-bound interfaces. The router is shown correctly in the atlas. If IP is not running on a third-party router but NXPIP is running on the management server, NetExplorer discovers the IPX-bound interfaces. However, these IPX-bound interfaces are not placed in the same router icon in the atlas.

On-Demand Links

An on-demand link is a WAN connection between two routers in which only user data (no routing traffic) is exchanged across the link. The link is brought up only when there is data to send.

NetExplorer discovers on-demand IP and IPX links correctly, if sufficient static routing information has been configured to allow the management server to reach the other side of the on-demand link.

However, if a link is an on-demand and unnumbered IP link, the entire topology on the remote end of the link is not discovered. Click IP Discovery, and then the Additional IP Routers in the NXPCON utility to configure an additional IP router address for the missing router.

Third-Party Routers with WAN Ports

NetExplorer discovers third-party routers correctly if they support MIB-II SNMP. Certain third-party routers can have a WAN link with no IP or IPX network number on the link. In this case, the WAN link is not discovered.

Novell NetWare Connect Servers

NetExplorer discovers Novell NetWare Connect servers; however, if you have more than one Novell NetWare Connect[®] server on the network, NetExplorer consolidates them and they appear as one server.

Virtual Switches

A virtual switch is represented by the same icon used for a switch or bridge in the atlas maps. The display name of a virtual switch is always shown as the “switch on *IP address of network*.” It is primarily used in atlas maps to display a meaningful network topology when discovery information is incomplete. Since the addresses are not known, the MAC address of the virtual switch is specified as an *Ethernet Port* or *Token Ring Port* or *FDDI Port* based on the connectivity type.

A virtual switch is shown in atlas maps under the following conditions:

- ◆ When two or more different physical media are connected by a switch, but the switch is not yet discovered. The virtual switch disappears as soon as the real switch is discovered.

- ◆ When two or more different physical media are connected by a switch, the switch is configured with SNMP community strings other than public, and the SNMP community strings of the switch were not provided through NXPCON before starting discovery.
- ◆ When two or more different physical media are connected by a non-manageable switch or a hub.

Network Segments

NetExplorer discovers the following network segments:

- ◆ “LAN and WAN Segment Types” on page 888
- ◆ “Source-Route Bridged Token Rings” on page 888

NetExplorer cannot fully discover the following:

- ◆ “Transparent Bridges” on page 889
- ◆ “Configuration Changes” on page 890

LAN and WAN Segment Types

NetExplorer discovers the LAN and WAN segment types shown in [Table 23-6](#):

Table 23-6 List of Known and Unknown segments in CIM database

| Known Segments in CIM Database | Unknown Segments in CIM Database |
|--------------------------------|----------------------------------|
| ATM | LAN: ARCnet |
| LAN: FDDI | LAN: LocalTalk* |
| LAN: Ethernet | SMDS |
| LAN: Token Ring | WAN: ISDN |
| WAN: X.25 | WAN: SDLC |
| WAN: PPP | WAN: Serial |
| WAN: Frame_Relay | WAN: T1 |
| | WAN: T3 |

These values are discovered correctly if a system connected to the segment responds with an interface type from MIB-II RFC 1573.

Source-Route Bridged Token Rings

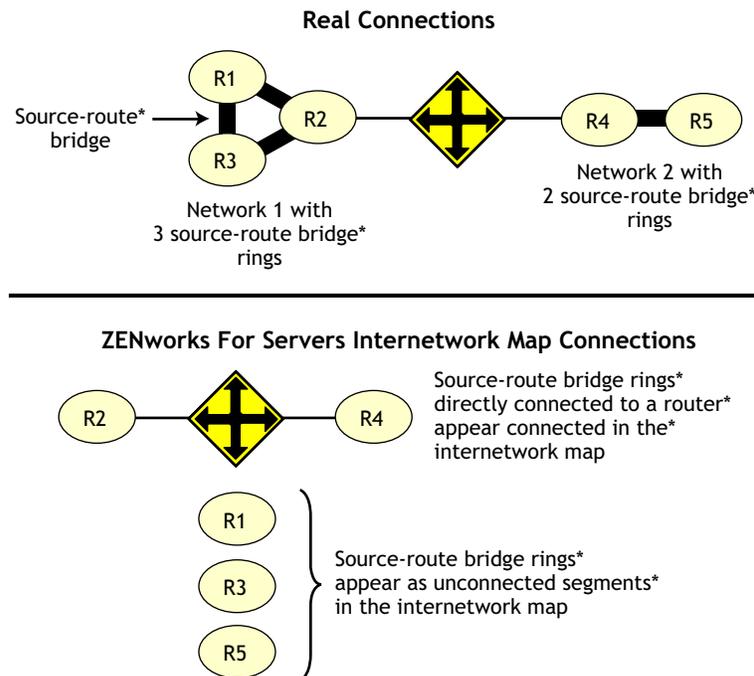
Atlas Manager displays source-route bridged token rings depending on whether the Traffic Analysis Agent for NetWare is installed on each ring.

- ◆ If you do not have the Traffic Analysis Agent for NetWare installed on each source-route bridged token ring in your network, NetExplorer discovers the network but consolidates all source-route bridged token rings that share the same IPX network number or IP subnet into a single segment. For example, in [Figure 23-14](#), rings R1, R2, and R3 are displayed as one segment, and rings R4 and R5 are displayed as another segment on the internetwork map.

- ◆ If you have the Traffic Analysis Agent for NetWare installed on each source-route bridged token ring, each Traffic Analysis Agent for NetWare discovers its own ring (segment) and every system on it. Atlas Manager displays the ring as a disconnected segment on the internetwork map.
- ◆ If you have the Traffic Analysis Agent for NetWare installed on a source-route bridged token ring connected to a router, the WAN page in the atlas shows the correct connections. However, if two networks each have several rings and only one ring in each network is connected to a router, the WAN page shows the correct connections of only the rings that are directly connected to the router. The other source-route bridged token rings in each network are displayed as disconnected segments on the WAN page.

Figure 23-14 illustrates this second case:

Figure 23-14 Internetwork Map Connections for Source Route Bridge Rings



In all cases, bridge information is not discovered. As a result, discovery treats each interface of a source-route bridge as a separate system on the network. One icon appears in the atlas for each interface of the source-route bridge.

When you have the Traffic Analysis Agent for NetWare installed on one server on each ring of an IPX source-route bridged network, the segment names displayed on the WAN page consist of the IPX network number followed by the MAC address of that server's interface to the ring. If the Traffic Analysis Agent for NetWare is monitoring more than one interface, the address shown for a ring is the MAC address of the interface monitoring that ring.

Transparent Bridges

Discovery cannot completely discover transparent bridges. It consolidates groups of transparently bridged segments running the same network number into a single segment on the maps.

Configuration Changes

Discovery detects most changes in the network topology, such as the addition, reconfiguration, or deletion of interfaces, resulting in changes being made to the atlas. However, if you remove the system from the network, it is not detected unless you move it to another location in the network.

23.1.4 File-Based Discovery

The enhancement to the `ipgproper.nlm` allows you to use the `discnodes.txt` file to specify the IP Address and mask of a set of nodes to be discovered. The information about the nodes is obtained through SNMP.

The IPGProper NLM must be loaded with specific options that enable it to receive inputs from the `discnodes.txt` file. If these options are not provided, the NLM will discover without taking input from the file. Prior to starting the discovery, the `discnodes.txt` file must be placed in the `ZENworks_installation_directory/mwserver/nmdisk` directory. After the initial discovery, if you want more nodes to be discovered, you must create a new `discnodes.txt` file with the new node entries and place it in the same directory. These nodes will be queried in the next discovery cycle.

The `discnodes.txt` input file has the following format for individual IP addresses:

- ◆ Individual IP Address specification format:

IPAddress <, SubnetMask>

- ◆ Specifying addresses using regular expressions

IPAddress <, SubnetMask>

IPAddress ->AddressPattern

Characters allowed in AddressPattern include the numerals 0-9; the period (.); the question mark (?), which represents one character; and the asterisk (*), which represents more than one character, up to a maximum of three.

- ◆ Wildcard characters are not allowed in the subnet mask.

| | |
|--------------|--|
| 164.99.149.* | All addresses in the range from 164.99.149.1 to 164.99.149.254 |
| 164.99.14?.* | all addresses in the range from 164.99.140.1 to 164.99.149.254 |
| 164.99.149.? | all addresses in the range from 164.99.149.1 to 164.99.149.9 |

NOTE: 164.99.149.0 does not come into the range. ? does not stand for 0 if it is the only letter in the octet.

164.99.149.1?0- all addresses in the range from 164.99.149.100 to 164.99.149.190. Here ? stands from 0

- ◆ In the text file, any line that begins with a pound sign (#) is treated as a comment line.

File-based discovery can be used in the following two scenarios:

- ◆ [“Discovering the Nodes Specified in the file” on page 891](#)
- ◆ [“Discovering the Nodes with Other Discovery Modules” on page 891](#)
- ◆ [“Using Command Line Options for IPGROPER” on page 892](#)

Discovering the Nodes Specified in the file

By default, the Novell ZENworks Server Management installation loads the NXPCON utility with all the discovery modules running and with file-based discovery enabled.

To discover only the nodes specified in the input file:

- 1 In NXPCON, click *Configuration Options > Discovery Modules*.
- 2 Select *Individual Discovery Modules*, then press Enter.
- 3 Select *No* to unload the modules, then press Enter.
- 4 Press Esc to exit the Discovery Modules dialog box.
- 5 Click *Yes* to save changes.
- 6 Click *Configuration Options > IP Discovery*.
- 7 Select *IP Host Discovery*, then press Enter.
- 8 Select *Enable IP Host Discovery*, then press Enter.
- 9 Select *No* to disable autodiscovery of the IP workstation.
- 10 Make sure that the Enable File-Based Discovery option is set to *Yes*.
- 11 Press Esc to exit the IP Host Discovery dialog box.
- 12 At the Management server prompt, unload NetExplorer by entering `unxp`.
- 13 Reload the NetExplorer modules by entering `netexplor`.

Discovering the Nodes with Other Discovery Modules

By default, the Novell ZENworks Server Management installation starts all the discovery modules along with the file-based discovery. Use the following procedure to individually select the modules that need to be started or to change the configuration.

To discover only the nodes specified in the input file:

- 1 In NXPCON, click *Configuration Options > Discovery Modules*.
- 2 Select *Individual Discovery Modules*, then press Enter.
- 3 Select *Yes* or *No* to load or unload each module, then press Enter.
- 4 Press Esc to exit the Discovery Modules dialog box.
- 5 Click *Yes* to save changes.
- 6 Click *Configuration Options > IP Discovery*.
- 7 Select *IP Host Discovery*, then press Enter.
- 8 Select *Enable IP Host Discovery*, then press Enter.
- 9 Select *No* to disable Auto Discovery of the IP workstation.
- 10 Make sure that the *Enable File-Based Discovery* option is set to *Yes*.
- 11 Press Esc to exit the IP Host Discovery dialog box.
- 12 At the Management server prompt, unload NetExplorer by entering `unxp`.
- 13 Re-load the NetExplorer modules by entering `netexplor`.

Using Command Line Options for IPGROPER

The `ipgroper.nlm` has three command line options to discover nodes specified in the `discnodes.txt` file.

Table 23-7 *Command Line Options for IPGROPER*

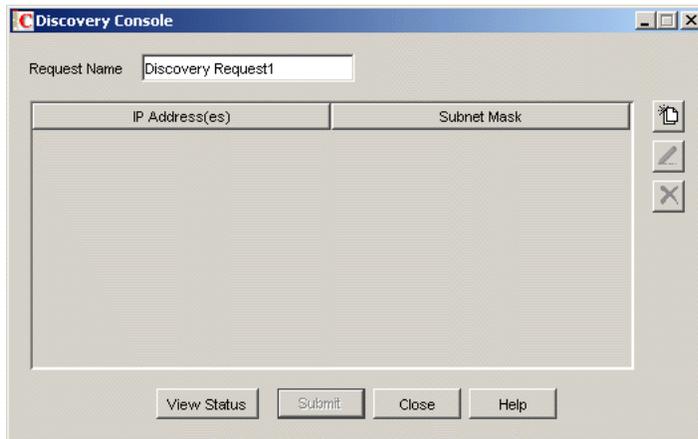
| Command Line | Explanation |
|---------------------|---|
| <code>/Fonly</code> | <p>Specifies that only the nodes specified in the <code>discnodes.txt</code> file must be discovered.</p> <p>You can set this option, when you have set <i>No</i> for the <i>Enable IP Host Discovery</i> option and <i>Yes</i> for the <i>Enable File-Based Discovery</i> option.</p> |
| <code>/Falso</code> | <p>Specifies that the nodes specified in the <code>discnodes.txt</code> file must be discovered along with the other nodes that IPGROPER will discover.</p> <p>You can set this option, when you have set <i>Yes</i> for both <i>Enable IP Host Discovery</i> option and <i>Enable File-Based Discovery</i> option.</p> |
| <code>/Flog</code> | <p>Logs all the errors and events that occur during the discovery of the nodes. The errors and the events will be logged in the <code>ZENworks_installation_directory\mwserver\nmdisk\discnodesbak\discnodeslog.log</code> file. You must manually enter this command line option in the <code>netxplor.ncf</code> file.</p> <p>IMPORTANT: The <code>discnodeslog.log</code> file will be created only if you have specified the <code>/fonly</code> or the <code>/flog</code></p> |

23.1.5 Discovery Console

The Discovery Console enables you to send a request to discover a set of IP addresses using Novell ConsoleOne. You can discover a list of host addresses, all the hosts on a subnet, range of addresses, or addresses in the form of a regular expression. The Discovery Console also enables you to view the status of the requests that you have submitted, or to delete a request.

IMPORTANT: To use the Discovery Console, ensure that IPGROPER is running.

Figure 23-15 *Discovery Console*



To create and submit a request for discovery:

- 1 From Novell ConsoleOne, click *Tools > Discovery Console*.

The Discovery Console dialog box is displayed with a default Request Name.

- 2 Enter a different request name if you want to change the default request name.

You can perform the following operations:

- ♦ “Add IP Addresses” on page 893
- ♦ “Edit IP Addresses” on page 894
- ♦ “Remove IP Addresses” on page 894
- ♦ “View Status of a Request” on page 894

Add IP Addresses

- 1 In the Discovery Console dialog box, click .



2 Provide the following information:

- ♦ **IP Address:** The host address or a subnet address. You can include wildcard characters while specifying the IP address. Specify the correct subnet mask for a subnet address.
- ♦ **IP Address Range:** The range of IP addresses to query. For example, 160.100.144.1 - 160.100.144.254.
- ♦ **Subnet Mask:** The correct subnet mask for a specified address.

3 Click *OK*.

4 In the Discovery Console dialog box, click *Submit*.

Edit IP Addresses

1 In the Discovery Console dialog box, select the IP addresses you want to edit, then click .

2 Modify the required information.

3 Click *OK*.

4 In the Discovery Console dialog box, click *Submit* to submit the request.

Remove IP Addresses

You can remove IP addresses if you have not added the request.

1 In the Discovery Console dialog box, select the IP addresses you want to remove, then click .

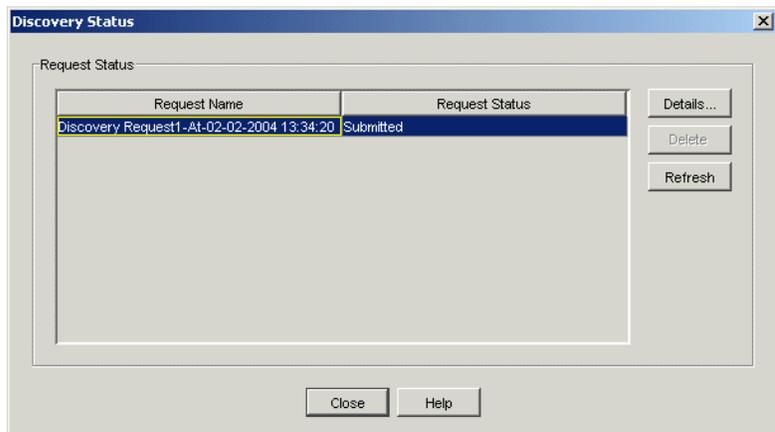
View Status of a Request

This option displays the status details of the request that you have submitted. A request can have one of the following status levels:

- ♦ **Pending:** The request is yet to be submitted to IPGROPER.
- ♦ **Submitted:** The request is submitted to IPGROPER.
- ♦ **In Progress:** The request is being processed by IPGROPER.
- ♦ **Completed:** The request is processed by IPGROPER.

To view the status of the request:

- 1 In the Discovery Console, select the request, then click *View Status*.



You can perform the following operations:

- ♦ **Details:** Displays the details for the selected request in the table.
- ♦ **Delete:** Deletes a request you selected. You can delete a request if the status is pending.
- ♦ **Refresh:** Displays the current status of the requests.

23.1.6 Effects of Discovery on Maps

The Atlas Manager on the management server creates an atlas database as the topology database is populated and the information is displayed as maps on Novell ConsoleOne. The WAN page displays all the Area pages and the connecting routers between them. The Area pages display the segments and the connecting routers.

The discovered systems are placed on the Area pages of the atlas based on the connecting routers or bridges. The Islands page contains segments for which routers have not yet been discovered. The Atlas Manager relocates the segments to the correct pages when connecting routers are discovered.

Review the following sections for more information on the effects of discovery on maps:

- ♦ [“Name Source Priority” on page 895](#)
- ♦ [“Representation of Systems in the Atlas” on page 896](#)

Name Source Priority

As discovery cycles proceed and more information is discovered, the names displayed in the maps can change. Different priorities are given to names, depending on the source of the name information. If none of the names are discovered, the IP/IPX address of the node is displayed as the node name.

To determine how to display the name of the discovered object, the Atlas Manager uses the following list in the order shown:

1. User Defined Name
2. DNS Name

3. Novell eDirectory Name
4. Bindery Name
5. SNMP Name

Representation of Systems in the Atlas

When representing a system in a map, the Atlas Manager refers to the following list of services in the order shown in [Table 23-8](#). As soon as it associates the first service with the node, it displays it without looking for further matches. The icon may change if a service with a higher priority is detected later during discovery.

Table 23-8 *Atlas Manager Services*

| Priority Number | Icon | Description |
|-----------------|---|--|
| 1. |  | NetWare server running the server management agent software |
| 2. |  | Windows server running the server management agent software |
| 3. |  | SFT III server running the MS engine |
| 4. |  | Server running file server software |
| 5. |  | Router running IP service |
| 6. |  | Router running IPX service |
| 7. |  | A switch or a bridge |
| 8. |  | Server running the discovery process |
| 9. |  | Server running the topology database |
| 10. |  | NetWare or Windows server running the traffic analysis (Traffic Analysis) agent software |

| Priority Number | Icon | Description |
|-----------------|---|---|
| 11. |  | Server running Remote Monitoring |
| 12. |  | Server running Remote Monitoring II |
| 13. |  | Server running print server software |
| 14. |  | Server running IP software |
| 15. |  | Server running NetWare Connect software |
| 16. |  | Router |
| 17. |  | Printer |
| 18. |  | IP workstation |
| 19. |  | IPX workstation |
| 20. |  | Others |

If a system has either an IPX or IP router service, the Atlas Manager considers it a router and displays it on the appropriate pages and segments.

23.2 Setting Up Discovery

The discovery software on a management server automatically discovers the nodes on your network. Network nodes include servers, desktops, routers, hubs, switches, and any other network devices. The Consolidator on the server populates the database with the discovered data. The Atlas Manager on the server reads the database and creates the atlas.

Novell ZENworks Server Management allows discovery in two different environments:

- ◆ Pure IP environment

- ◆ IP/IPX environment

You must have IP enabled between Novell ConsoleOne and the management server.

Before starting discovery, you must verify the following configurations to ensure that the discovery system is complete:

- ◆ Ensure that the router to which Novell ZENworks Server Management server is attached is specified as the seed router in NXPCON. If necessary specify, additional IP routers also. For more information on specifying seed router and additional IP routers, see [“Specifying a Seed Router and Additional IP Routers” on page 906](#)
- ◆ Ensure that the community strings used for all the devices to be managed are specified in NXPCON. For more information on changing SNMP community strings, see [“Changing the SNMP Community String” on page 902](#).
- ◆ Ensure that the Novell ZENworks Server Management server is privileged to query the routers in your network if the routers are configured to restrict access to only specified IP addresses. For more information on IP router discovery, see [“IP router discovery on IP networks only.” on page 874](#).
- ◆ If you want to restrict the scope of IP or IPX discovery, specify proper scoping entries. For more information on changing the discovery scope, see [“Changing the Discovery Scope” on page 903](#).
- ◆ Ensure that the DNS configuration file, `sys:\etc\resolv.cfg`, has a valid DNS server's IP address. If a valid DNS server is not specified, discovery will fail to discover the DNS names of hosts.
- ◆ For effective discovery, ensure that the Traffic Analysis Agent is installed and running on each network segment that you want to discover and manage. Also, ensure that the names and addresses of these agents are specified in NXPCON. For more information on specifying Traffic Analysis Agents, see [“Specifying Traffic Analysis Agents to Be Queried by NXPLANZ” on page 905](#).
- ◆ If a MAC address is being associated with different network numbers, all such network numbers will be merged into a single segment. To avoid the merger, you must specify all such MAC addresses in upper case in the `installation_directory\mms\mwserver\bin\consolidator.ini` file.

In `consolidator.ini`, specify the MAC address as a key value pair in the `[DuplicateMacAddress]` section.

A sample `consolidator.ini` is as follows:

```
[DuplicateMacAddress]
mac1="00C04F59910D"
mac2="00C04F5991AB"
...
key_name=value
```

In `consolidator.ini`, ensure that the keys are unique.

Before starting the Novell ZENworks Management and Monitoring Services server, edit the `ZENworks_installation_directory\mms\mwserver\properties\sloder.properties` file to append the ARGUMENTS value under TOPOLOGY MANAGER with the following entry:

```
-ini "installation_directory\mms\mwserver\bin\consolidator.ini"
```

The following tasks will start discovery initially and help you customize discovery to meet your organization's needs:

- ◆ [Section 23.2.1, “Starting Discovery,” on page 899](#)
- ◆ [Section 23.2.2, “Checking the Status of Initial Discovery,” on page 899](#)
- ◆ [Section 23.2.3, “Checking the Results of Discovery,” on page 900](#)
- ◆ [Section 23.2.4, “Changing the Default Configuration,” on page 901](#)
- ◆ [Section 23.2.5, “Configuring the Java Processes,” on page 906](#)
- ◆ [Section 23.2.6, “Unloading the Management Server,” on page 908](#)

23.2.1 Starting Discovery

Discovery starts automatically when the discovery software is loaded on the management server.

To manually start autodiscovery and load the back-end services (management site services), refer to the steps in [“Management and Monitoring Services Installation”](#) in the *Novell ZENworks 7 Server Management Installation Guide*.

Restarting the Management Server

If you bring down the management server (for example, for maintenance), the restart affects discovery in the following ways:

- ◆ Each time you reload the discovery modules, a new version of `netexplor.dat` is created.
- ◆ The initial discovery cycle starts again.
- ◆ The Consolidator processes all the discovery data again as Novell ZENworks Server Management rediscovers the network.

To unload the discovery modules:

- 1 At the NetExplorer server, enter `unxp`.

To load the discovery modules:

- 1 At the NetExplorer server, enter `netexplor`.

23.2.2 Checking the Status of Initial Discovery

As discovery progresses, your topology maps in Novell ConsoleOne reflect the discovered data. However, in a large network, it might take a day or two before the initial discovery is complete.

The easiest way to determine whether initial discovery is complete is to use the `NXPCON` utility on the management server and check the status of each NetExplorer module. Each module must complete at least one full cycle to draw a complete map.

To view the discovery status, look at the discovery status fields at the top of the `NXPCON` screen. See [“Using the Discovery Configuration Utility” on page 901](#) for information about how to access this screen.

The `NXPCON` main screen gives you the information you can use to monitor the status of discovery.

The following information is displayed:

- ◆ **NetExplorer Up Time:** Shows the time since NetExplorer started running.
- ◆ **NetExplorer System Status:** Shows the overall status. It can have one of the following values:
 - ◆ Waiting to start - Waiting for one or more of the discovery modules to start.
 - ◆ Running - Discovery modules are running.
- ◆ **Module Status:** Shows the status of each module and the number of cycles each module has completed. The module status can be one of the following values:
 - ◆ Not Loaded - Module is not loaded.
 - ◆ Waiting to Start - Module is loaded but not started.
 - ◆ Running - Module is running and collecting data.
 - ◆ Suspended - Module is suspended because it reached the end of the schedule in which it was running.
 - ◆ Completed - Module completed a discovery cycle.
 - ◆ Unknown - NetExplorer cannot obtain the module status. (This is usually seen if the module is not loaded.)

23.2.3 Checking the Results of Discovery

When the Consolidator has finished updating the database after the initial discovery, verify if the network topology is accurately represented on the maps.

NetExplorer might not have discovered the type if a node is not on the map. If a node does not appear in the correct segment, NetExplorer may not have received sufficient information to place it correctly. For more information, see [“What Is Discovered” on page 881](#). The following characteristics are captured:

- ◆ IP - Discovers IP routers; IP hosts; IP services such as HTTP, Telnet, SMTP, DNS, FTP; and DHCP.
- ◆ IPX - Discovers IPX workstations, IPX routers, and IPX services (file, print, any other Service Advertising Protocol [SAP]).
- ◆ Subnet mask
- ◆ Services
- ◆ Novell eDirectory names and tree
- ◆ DNS Names

The Consolidator on the management server communicates with NetExplorer to obtain network discovery data. The Consolidator reads the `netexplor.dat` file and populates the database.

IMPORTANT: The `netexplor.dat` file is reset every time you restart NetExplorer.

The Consolidator communicates with two Java* components: the Bridge Agent and the SN3 agent. The Bridge agent retrieves bridges present in the network and the related topology of the network. The SN3 agent does SLP-based discovery for NetWare 5.x servers and gets the corresponding Novell eDirectory name for each IP and IPX address discovered.

IMPORTANT: NetExplorer and the Consolidator can run independent of each other on the management server.

NetWare 5.x/6.x servers are discovered faster because NetWare 5.x/6.x supports the Service Location Protocol (SLP).

Ensuring Complete Discovery

IPX workstations are discovered with a username if the user is logged in to or attached to a NetWare server running management agent software. To ensure that the usernames for IPX devices and workstations on your network can be discovered, install a management agent on all NetWare servers where users log in.

If you want NetExplorer to discover AppleTalk* devices, you need to install the NetWare Traffic Analysis Agent on one server on each segment.

23.2.4 Changing the Default Configuration

The discovery software is installed with default configuration designed to work in most environments. However, if your network or the data on your database is not discovered, you need to reconfigure discovery.

Read the following sections for more information:

- ◆ [“Using the Discovery Configuration Utility” on page 901](#)
- ◆ [“Choosing Which Discovery Modules to Load” on page 902](#)
- ◆ [“Changing the SNMP Community String” on page 902](#)
- ◆ [“Changing the Discovery Scope” on page 903](#)
- ◆ [“Specifying Traffic Analysis Agents to Be Queried by NXPLANZ” on page 905](#)
- ◆ [“Specifying a Seed Router and Additional IP Routers” on page 906](#)
- ◆ [“Refreshing the SNMP Configuration Settings of NetExplorer Using Activate Changes” on page 906](#)

Using the Discovery Configuration Utility

You can use the NXPCON utility on the management server to change the discovery configuration. For example, you can change the scope of discovery or view the status of the initial discovery process.

To access the NXPCON utility:

- 1 Access the server console on the management server either directly from the server prompt or remotely.
- 2 If the discovery modules are already loaded on the server, click the *NetExplorer Console Utility* option in the Available Screens window.

or

If the discovery modules are not loaded, enter `netexplor` at the server prompt.

NXPCON is loaded automatically when NetExplorer is loaded and is accessible at the management server.

If NXPCON is not loaded on your management server, check to see if NetExplorer is running. If NetExplorer is running, enter `load nxpcon` at the system console prompt. If NetExplorer is not running, enter `netexplor` at the system console prompt.

Choosing Which Discovery Modules to Load

By default, the Novell ZENworks Server Management installation loads the NXPCON utility with all modules running. If you are not using IPX on your network, you can configure NXPCON to not load the NXIPX module.

IMPORTANT: Make sure TCP/IP is bound to at least one of your server's network boards.

To view or modify which modules are being loaded:

- 1 In NXPCON, click *Configuration Options > NetExplorer Modules*.
- 2 Select the field you want to change, then press Enter.
- 3 Select *Yes* or *No* to load or unload the module, then press Enter.
- 4 Press Esc to exit the NetExplorer Modules dialog box.
- 5 Click *Yes*.

You can enable IP host discovery or file-based discovery. To enable or disable:

- 5a Select *Configuration Options > IP Discovery*.
- 5b Select *IP Host Discovery* or *File Based Discovery*, then press Enter, then *Yes* to enable or *No* to disable the discovery option.
- 6 At the management server prompt, unload NetExplorer by entering `unxp`.
- 7 Reload the NetExplorer modules by entering `netexplor`.

Changing the SNMP Community String

In Novell ZENworks Server Management, the default community string is PUBLIC. If your organization's SNMP community string is not PUBLIC, reconfigure the SNMP community string in NXPCON.

NOTE: In order to prevent burdening the routers, some organizations add one more level of control by allowing only certain IP addresses to do SNMP queries to the routers. If this is true in your organization, make sure that the IP address given to the Novell ZENworks Server Management sever is privileged to query the routers in the network. Otherwise, the discovery will not be complete and incomplete network information will appear under "Islands" in the atlas.

To view, add, modify, or delete SNMP configuration information, such as community strings used for IP and IPX discovery:

- 1 In NXPCON, click *Configuration Options > SNMP*.
- 2 In the SNMP dialog box, click *Edit Community Name List*.
- 3 To add a community string, press Insert.

or

To modify a community string, click the community string, then press Enter.

or

To delete a community string, click the community string, then press Delete.

- 4 Press Esc, then click *Activate Changes* from the Configuration Options window. For more information about Activate Changes, see [“Refreshing the SNMP Configuration Settings of NetExplorer Using Activate Changes” on page 906](#).
- 5 Respond to the prompts accordingly.

For information about other configuration options in the SNMP window, see [“Using the Discovery Configuration Utility” on page 901](#), or Novell ConsoleOne online help.

Changing the Discovery Scope

By default, NXPCON is set to discover all IPX and IP networks. You can, however, limit the discovery scope.

You could, for example, limit discovery to discover the IPX addresses or the IP subnet addresses. If you are managing a large network, by setting the scope of discovery, you will be limiting the discovery to a section of your network, which will reduce the network traffic and in turn make your atlas more manageable

If you do not accurately specify the scope of discovery, you will not be able to discover your target device. Therefore it is imperative to specify in the scope, all the devices that are present in the path leading to the target device you want to discover.

For example, consider the following scenario.

Your discovery server D1 is connected to network N1. Router R1 connects network N2 with N1. Assume you need to discover network N2. To do this, the following entries need to be set in the scope:

- ◆ Discovery server D1 with subnet mask 255.255.255.255
- ◆ Router R1 with subnet mask 255.255.255.255
- ◆ Network N2 with its appropriate subnet mask number.

In this scenario, network N2 can be reached from the discovery server through Router R1, and therefore R1 needs to be in the scope even if the user is not interested in the network N1 that R1 is routing.

After initial discovery, until you reset the database, nodes remain in the database even if they have been removed from the network.

Changing the discovery scope does not affect devices that are already in the database due to prior runs of discovery. In particular, devices that were discovered due to a wider scope (or no scope) will not be removed when a restrictive scope is set for later runs of discovery. If it is desired that the atlas shows only those devices that fall in scope, the database needs to be reset to ensure that segments and devices that are out of scope do not appear in atlas. Note that the database being reset would result in loss of data like alarms and alarm disposition unless they are migrated. Alternatively, if the number of such devices which are out of scope is very small, the user can manually delete them from the database using the Database Object Editor.

You can restrict the scope of IP or IPX discovery by entering the IPX network numbers or IP address ranges specified by the mask fields you want to discover. To view or restrict the IP or IPX scope:

- 1 In NXPCON, from the Configuration Options window, click *Discovery Scope*.

- 2 Select *IP Discovery Scope* or *IPX Discovery Scope*.
- 3 Press Enter to view or configure the scope of your discovery.
- 4 Press Insert to add a new IP or IPX discovery scope entry.

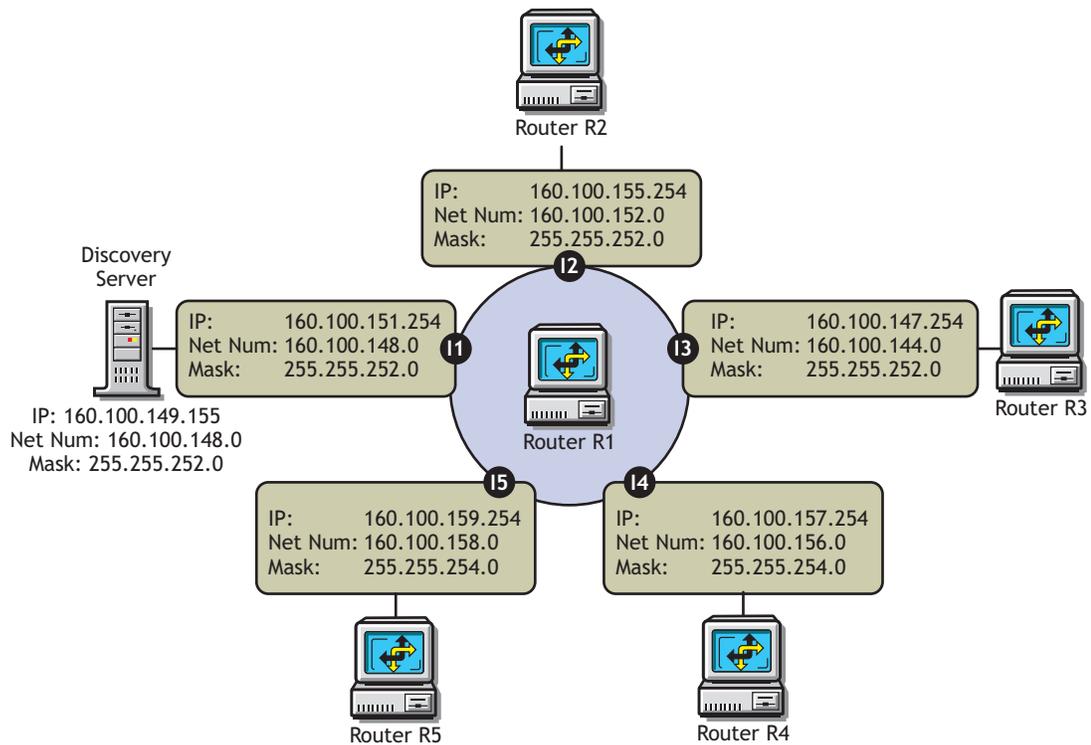
or

Press Enter to modify a discovery scope entry.

or

Press Delete to delete a discovery scope entry.

For IP Networks: Discovery scope is tightly bound to the network numbers. The scope can be restricted by specific networks as illustrated in the following diagram.



Case 1: To exclude 160.100.148.0 and discover the other four networks, specify the scope as:

160.100.149.155,255.255.255.255

160.100.151.254,255.255.255.255

160.100.144.0,255.255.252.0

160.100.152.0,255.255.252.0

160.100.156.0,255.255.254.0

160.100.158.0,255.255.254.0

The 255.255.255.255 mask for the Novell ZENworks Server Management server and the router interface on the local network acts as a machine specific scope. This prevents other machines in the network 160.100.148.0 from being discovered.

Case 2: To discover only the local network 160.100.148.0, specify the scope as:

160.100.148.0,255.255.252.0

The network 160.100.148.0 (Mask: 255.255.252.0) has IP addresses in the range 160.100.148.1 to 160.100.151.254.

Consider a case where all the important servers in your network have IP addresses in the range 160.100.149.1 to 160.100.149.254. You might specify the following scope:

160.100.149.0,255.255.255.0

The above scope is not allowed by the discovery system. You cannot set a scope to discover only a part of the subnet. You will have to set the entire subnet in scope.

Case 3: To discover only 160.100.156.0 and 160.100.158.0 scope should be given as:

160.100.149.155,255.255.255.255

160.100.151.254,255.255.255.255

160.100.156.0,255.255.254.0

160.100.158.0,255.255.254.0

Replacing the last two scoping entries with a single entry 160.100.156.0,255.255.252.0 might not have the same effect.

You cannot create a single scoping entry to cover two or more subnets. You have to create a scope for each subnet.

For IPX Networks: Restrict the scope to the IPX networks to be discovered by entering a single IPX network number and a mask.

The mask indicates which part of the network number needs to match. An F in the mask means that the corresponding digit must match; a 0 (zero) means that no match is required.

For example, network number 12340000 and mask FFFF0000 will match any network number starting with 1234.

Network number C00000FF and mask FF0000FF will match any network number starting with C0 and ending with FF, such as C01234FF or C00000FF.

- 5 Enter the address and mask for your discovery scope.
- 6 Press Esc, then click *Yes* to save changes to the configuration file.
- 7 Press Esc to return to the Discovery Scope window.
- 8 Unload and reload the NetExplorer modules or restart your management server for the changes to take effect.

Specifying Traffic Analysis Agents to Be Queried by NXPLANZ

Traffic analysis agents in your network are usually discovered by the NXPLANZ module. If SLP is disabled or if SAP packets are filtered by the routers in your network, NXPLANZ might not be able to discover all the Traffic Analysis Agents in the network.

To specify Traffic Analysis Agents to be queried by the NXPLANZ module:

- 1 In NXPCON, click *Configuration Options > NXPLANZ Discovery*.
- 2 To add an agent, press Insert.
- 3 Enter the address and mask for your discovery scope.
- 4 Press Esc, then click *Yes* to save changes to the configuration file.
- 5 Unload and reload the NetExplorer modules or restart your management server.
- 6 To modify an agent, select the agent, then press Enter. Modify the required information.

- 7 To delete an agent, select the agent, then press Delete.

Specifying a Seed Router and Additional IP Routers

Seed router is the router to which Novell ZENworks Server Management server is connected. For router discovery to be effective, always specify the seed router using NXPCON and ensure that Novell ZENworks Server Management server can query the seed router by specifying the proper community name in NXPCON.

You need to specify additional IP routers if you want to discover one part of your network and the Novell ZENworks Server Management server does not have access to one of the intermediate routers.

To specify a seed router or additional IP Routers:

- 1 In NXPCON, click *Configuration Options > IP Discovery > IP Router Discovery*.

The default for IP Seed Router is `<local>`, which is the Novell ZENworks Server Management server.

- 2 To add a seed router, select *IP Seed Router* and press Enter.
- 3 Enter the IP address.
- 4 To add additional routers, select *Additional IP Routers* and press Enter.
- 5 Enter the IP address.
- 6 Press Esc, then click *Yes* to save changes to the configuration file.
- 7 Unload and reload the NetExplorer modules or restart your management server.

Refreshing the SNMP Configuration Settings of NetExplorer Using Activate Changes

When you change the SNMP configuration settings of NetExplorer (such as the SNMP community string), the NetExplorer can be automatically updated with the changes without restarting NetExplorer. To automatically update NetExplorer with new settings, select *Activate Changes* from the Configuration Options window of NXPCON.

However, if you change any other configuration options of NXPCON (for example, the Discovery Scope), you must restart NetExplorer for these changes to be applied on NetExplorer. For more information on how to restart NetExplorer, see [“Stopping and Starting the Discovery NLM Files” on page 907](#).

23.2.5 Configuring the Java Processes

The following are the three Java processes of the discovery system:

- ◆ Topology Manager
- ◆ Bridge Discovery
- ◆ SN3 Discovery

These Java processes form a part of the Management Site Server and exist as sections in the `loader.properties` file in the `installation_path\novell zenworks\mms\mwserver\properties` directory. They are specified in the following format:

```
[Topology Manager]
Name = Topology Manager
Load Option = auto
Other options
```

To configure the Java processes:

1. Change the value of the Load Option from Auto to Manual to prevent the process from starting the next time you enter the SLOADER command on the server.

IMPORTANT: If you modify the `sloader.properties` file after you start the Management Site Server, you must restart the Management Site Server for the changes to take effect.

2. Do not change the Load Properties and the Load Sequence options in the `sloader.properties` file. These options are necessary for the Management Site Server to work correctly.

Customizing Starting and Stopping Discovery

You can choose to stop or start the discovery NLM files or the Java discovery processes without affecting the other services of the site server, such as the Alarm Manager Service.

Stopping and Starting the Discovery NLM Files

To stop the discovery NLM files, enter `UNXP` at the server console.

To start all the discovery NLM files, enter `NETXPLORE` at the server console.

NOTE: You cannot start the Discovery NLM files if the Java processes are running. Stop the Java processes and then enter `NETXPLORE` at the server console to start all the discovery NLM files.

Stopping and Starting the Java Discovery Services

To stop the discovery NLM files, enter `STOPDIS` at the server console.

To start all the discovery NLM files, enter `STARTDIS` at the server console.

You can customize starting or stopping any of the Java discovery processes at any point in time. For example, you decided not to run the Bridge discovery initially but decide to run it anyway. In such a scenario, you need not stop all the services and restart them.

You can edit the `startdis.ncf` file in the `\Novell ZENworks\mms\mwserver\bin` directory, which has the following contents:

```
MWSETENV.NCF
java -Xbootclasspath/p:$mwxbpath -classpath
$MMSCP;$CLASSPATHcom.novell.utility.servicemanager.ui.Start "Topology
Manager" "Bridge Discovery" "SN3 Discovery" <ip address of the server>
sloader
```

In the above file, the Java discovery process names like SN3 Discovery must match the names of the sections in the `sloader.properties` file. By changing just the names in the NCF files, you can create similar NCF files to selectively stop and start the Java discovery services.

For example, if you want to start just the Bridge discovery process:

1. Create a `startbri.ncf` file with the following contents:

```
mwsetenv.ncf
java -Xbootclasspath/p:$mwxbpath -classpath
$MMSCP;$CLASSPATHcom.novell.utility.servicemanager.ui.Start
"Bridge Discovery" <ip address of the server> sloder
```
2. Copy the `startbri.ncf` file to the `\novell zenworks\mms\mwserver\bin` directory.
3. Run the `startbri.ncf` file to start the Java discovery bridge service.

For example, to stop the Java discovery process for the SN3 Agent:

1. Create a `stopsn3.ncf` file with the following contents:

```
mwsetenv.ncf
java -Xbootclasspath/p:$mwxbpath -classpath
$MMSCP;$CLASSPATHcom.novell.utility.servicemanager.ui.Stop "SN3
Discovery" <ip address of the server> sloder.
```

23.2.6 Unloading the Management Server

To unload the management server:

- 1 If restarting the server is not feasible, make sure all Service Loader processes are exited. At the server console prompt, enter

```
stopService.ncf
```

This process might take some time to unload the services.

To know the status of “`com.novell.zfs.mms.utility.servicemanager.zfs_MMS_Services`” and “`com.novell.zfs.mms.utility.servicemanager.zfs_MMS_LanzSlpDis,`” enter `java - show` at the server console prompt.

You can use `java -exit` if you can terminate all other Java processes. Unload Java, if all the services are not closed.

- 2 If you are not running any other components of ZENworks 7 Server Management or ZENworks 7 Desktop Management that use the naming server, enter `stopmms -n` to unload all ZENworks 7 Server Management components.
- 3 Switch to the Sybase* process by pressing `Ctrl+Esc` then enter `q` to terminate the Sybase database engine.

23.3 Managing the Atlas

After the initial discovery, you can stop discovery running on the management server. You can, however, continue to access the database through the Atlas Manager. The discovery cycle starts again the next time NetExplorer is up. The Consolidator populates the database and the Atlas Manager automatically updates the atlas pages.

Depending on the size of your network, writing data from the initial discovery cycle can take few minutes to several days. Subsequent discovery updates to the database require substantially less time.

- ◆ [Section 23.3.1, “Using the Atlas,” on page 909](#)
- ◆ [Section 23.3.2, “VLAN Atlas,” on page 915](#)
- ◆ [Section 23.3.3, “Using Unified Views,” on page 916](#)

23.3.1 Using the Atlas

When Novell ZENworks Server Management is first installed, the server module of the Atlas Manager is automatically installed on the management server, and the client module of the Atlas Manager is installed on Novell ConsoleOne. The Atlas Manager on the management server creates a system atlas and provides a graphical view of the database at the console.

The Atlas Manager on the server reads the database and provides two different views of the database at Novell ConsoleOne: the Console view and Atlas view. Both views provide information about the discovered network topology, the physical location of nodes, node configuration information, and alarm information.

The following sections gives you an understanding about using the atlas:

- ◆ [“Accessing the Atlas” on page 909](#)
- ◆ [“Assigning Roles to Help You Manage the Atlas” on page 910](#)
- ◆ [“Using the Atlas to Troubleshoot” on page 910](#)
- ◆ [“Custom Maps” on page 911](#)
- ◆ [“Node Naming Order” on page 915](#)

Accessing the Atlas

You can access the Novell ZENworks Server Management atlas from Novell ConsoleOne. Open Novell ConsoleOne and double-click the Novell ZENworks Server Management Domains namespace, then expand the domain. The system atlas appears.

[Table 23-9](#) describes a Novell ZENworks Server Management atlas consisting of three different pages:

Table 23-9 *Three-page Atlas*

| Atlas Pages | Icon | Description |
|-------------|---|--|
| WAN page |  | Summarizes the entire network, illustrating the WAN-related network topology. Your atlas, typically, has a single WAN page. |
| Area page |  | Displays segments on your network. An atlas can have several area pages. For example, areas can be divided based on the geographic location of the network. If a company in San Jose has an overseas branch in Germany, you can divide your network into Area1 for the San Jose network and Area2 for the Germany network. |

| Atlas Pages | Icon | Description |
|--------------|---|---|
| Islands page |  | Consists of segments with an undetermined connectivity. During discovery, the Islands page is a placeholder for network objects that are not completely discovered. An atlas has a single Islands page. |

Customizing Your Atlas View

You can customize your atlas view in three different ways:

- ◆ Insert a custom bitmap as the background on an atlas page.
- ◆ Change the position of a node on an atlas page by dragging it.
- ◆ Display objects by an alternate name.

Assigning Roles to Help You Manage the Atlas

Novell ZENworks Server Management lets you assign roles to manage the atlas. By assigning roles, you can restrict the user from performing specific operations on that object.

TIP: The atlas displays maps based on your role on the network. For example, if your role is restricted to managing certain servers in segment A and B, your atlas will contain only those servers in segments A and B.

You can perform the following tasks on any atlas page (WAN, Islands, or Area page) when the Atlas view is displayed on Novell ConsoleOne:

Table 23-10 Tasks that can be performed on an Atlas page

| Tasks | Comments |
|--------|--|
| Open | Opens the page |
| Import | Inserts a custom wallpaper |
| Save | Updates the changes in the database |
| Print | Prints the page |
| Rename | Renames the page |
| Layout | Displays the page with a different focal point |

Using the Atlas to Troubleshoot

By setting the alarm disposition to save alarms in the database, Novell ZENworks Server Management maps can alert you to alarm conditions on the network. Alarms are of type severe, major, or minor alarm on a segment or node. Upon recognizing any of these alarms, the ConsoleOne displays different colors above the object depending on the severity of the alarm as shown in [Table 23-11](#):

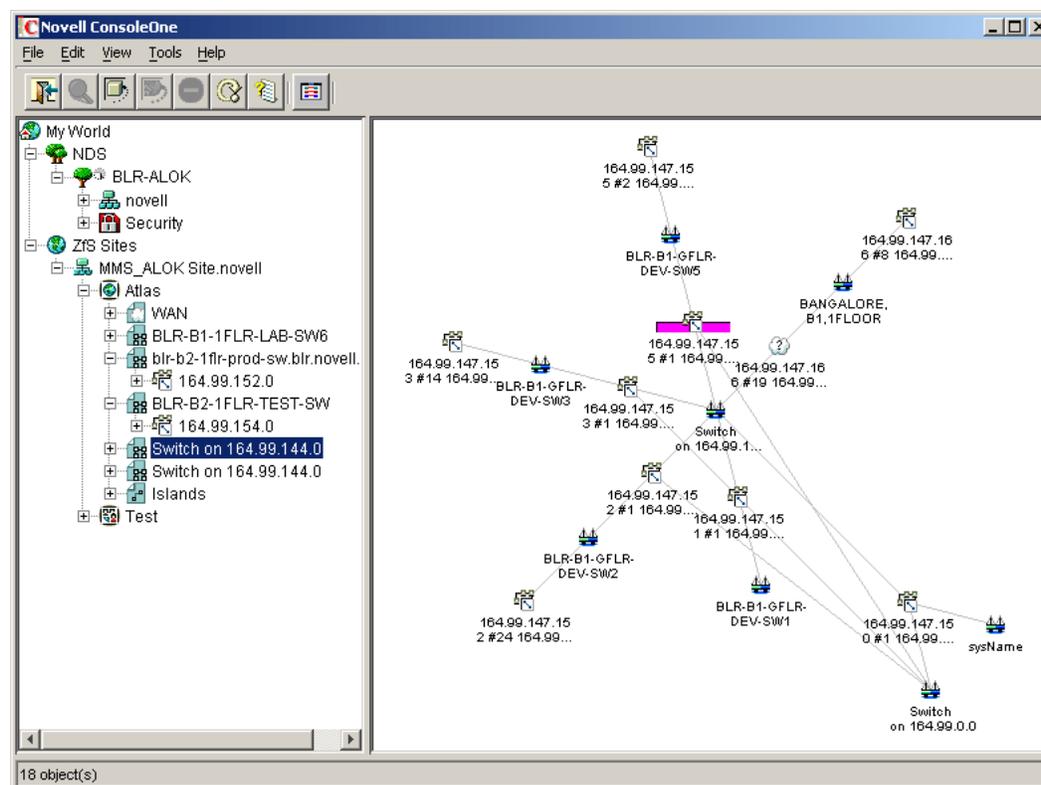
Table 23-11 Alarm Severity

| Color | Severity of Alarm |
|--------|-------------------|
| Red | Sever |
| Pink | Major |
| Yellow | Minor |

The alarm status is propagated up the hierarchy. For example, if a server has an alarm of type severe, the segment and the page containing the server will display the corresponding alarm icon. For information about alarms, [Section 24.2, “Managing the Alarm Management System,” on page 924](#).

[Figure 23-16](#) shows the atlas namespace in ConsoleOne:

Figure 23-16 Atlas namespace in Novell ConsoleOne



Custom Maps

Custom Maps enable you to create and delete custom atlases and custom containers and group nodes into containers. You can also create a hierarchy of objects in atlas. An atlas can contain custom containers. you can create a node or a sub container within the custom container. However nodes cannot be directly contained under the atlas.

You can perform the following operations in the custom atlas:

- [“Creating a Custom Atlas” on page 912](#)
- [“Creating a Custom Container” on page 912](#)

- ◆ “Adding Nodes to the Container” on page 913
- ◆ “Renaming the Custom Atlas Objects” on page 913
- ◆ “Deleting Custom Atlas Objects” on page 913
- ◆ “Locating Nodes in the Custom Atlas” on page 913
- ◆ “Locating All Occurrences of the Nodes” on page 914
- ◆ “Copying Nodes to a Custom Container” on page 914

Creating a Custom Atlas

- 1** Right-click the site object where you want to create the custom atlas, then click *New > Atlas*.
- 2** Specify the name of the atlas.
- 3** Click *OK*.

Creating a Custom Container

You can create a custom container within a new custom atlas. You can also create a subcontainer within a container.

- 1** Right-click the custom atlas you created, then click *New > Container*.

or

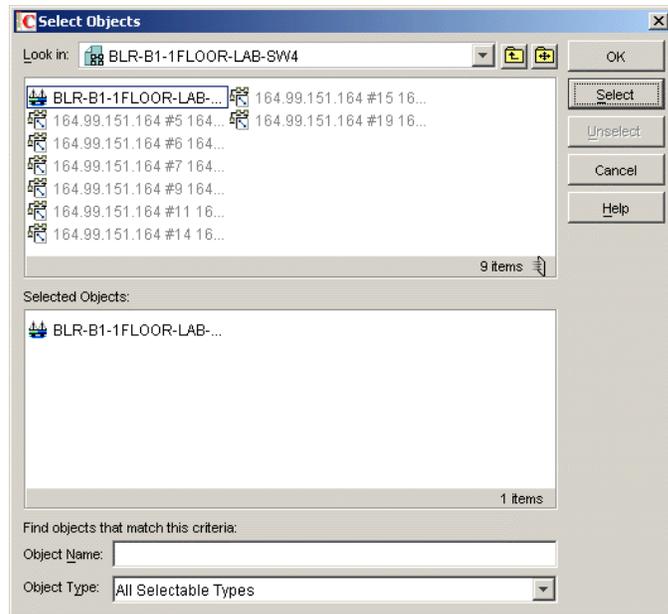
In the right-pane of Novell ConsoleOne, right-click the custom container you created, then click *New > Container*.

- 2** Specify the name of the container.
- 3** Click *OK*.

Adding Nodes to the Container

You can associate multiple nodes to a container.

- 1 Right-click the custom container you created, then click *Add Nodes*.



- 2 Select the nodes you want to add to the container.
- 3 Click *Select*.
The selected nodes are displayed in the *Selected Objects* list.
- 4 Click *OK*.

Renaming the Custom Atlas Objects

- 1 Select the atlas object you want to rename.
- 2 Click *Rename*.
- 3 Specify a new name.
- 4 Click *OK*.

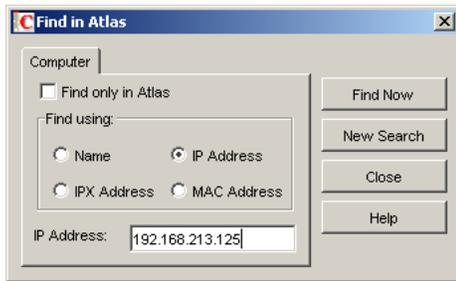
Deleting Custom Atlas Objects

- 1 Select the atlas object you want to delete.
- 2 Click *Delete*.

Locating Nodes in the Custom Atlas

You can locate particular node in the specific atlas or across all the atlases that are present. To locate nodes in the same atlas, select the Find only in *atlasname* check box. If you do not select this check box, then the node will be located in all the atlases.

Figure 23-17 Find in Atlas Dialog Box



- 1 Right-click the custom atlas, click *Find*.
- 2 Specify the conditions to base your search on.
- 3 Click *Find Now*.

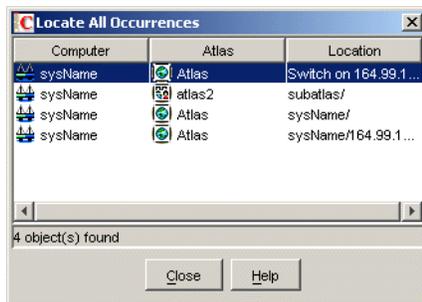
The location of the network object in the custom atlas is displayed in a list in the Find In Atlas dialog box.

- 4 To jump to the location of the node, double-click the node.
- 5 Click *Close*.

Locating All Occurrences of the Nodes

- 1 Right-click the node whose occurrences you want to locate, then click *Locate All Occurrences*.

The occurrences of the nodes are displayed in the Locate All Occurrences dialog box.



- 2 To jump to the location where the node exists, double-click the node.
- 3 Click *Close*.

Copying Nodes to a Custom Container

You can copy nodes to the custom container you have created.

- 1 Select the nodes you want to copy, then right-click and select *Copy to Container*.
- 2 In the Select Objects dialog box, locate the custom container where you want to copy the nodes to.
- 3 Click *OK*.

Node Naming Order

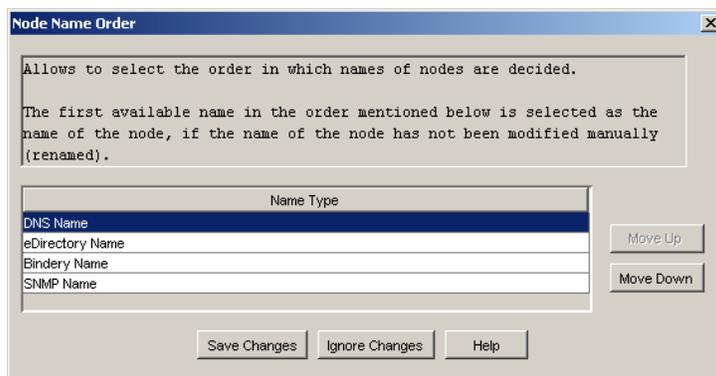
As discovery cycles proceed and more information is discovered, the names displayed in the maps can change. Different priorities are given to names, depending on the source of the name information. If none of the names are discovered, the IP/IPX address of the node is displayed as the node name.

To determine how to display the name of the discovered object, the Atlas Manager uses the following list in the order shown:

- ◆ User - defined name
- ◆ DNS Name
- ◆ Novell eDirectory Name
- ◆ Bindery Name
- ◆ SNMP Name

Use the Node Name Order dialog box to change the order of display.

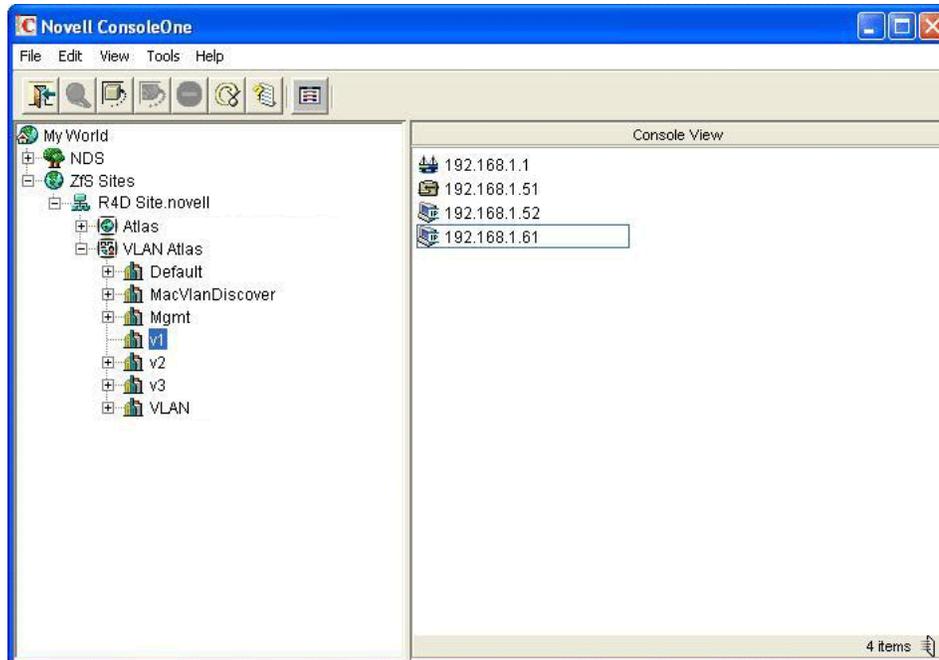
- 1 At the atlas level, select *File > Node Naming Order*.



- 2 Use the following options:
 - ◆ **Move Up:** Changes the order of the node names.
 - ◆ **Move Down:** Changes the order of the node names.
 - ◆ **Save Changes:** Saves the changes you made.
 - ◆ **Ignore Changes:** Ignores the changes you made.

23.3.2 VLAN Atlas

VLAN Atlas provides a logical view of all VLANs present in the network. VLAN feature uses the custom atlas framework to create a VLAN-specific atlas.



Under VLAN atlas, you can see all the networks running on different VLANs.

The naming convention for the VLANs is VLAN_name. All the machines belonging to a particular VLAN, are placed in respective containers. The assumption made is that, in a given organization the names of VLANs are distinct.

Because VLAN Atlas is a custom atlas, all the operations that are possible in a custom atlas are possible for VLAN Atlas as well. However, renaming or deleting a container is not possible.

Only port-based VLANs are currently supported on Baystack, Extreme, and Cisco switches.

23.3.3 Using Unified Views

The Unified view service is a service that acts as a filter on the atlas. Using the Unified view, you can filter for a list of devices or segments of a particular type. The Unified view allows easy navigation and quick operations to check the highest severity of the alarms present on a particular node or segment.

The following are the two types of Unified view provided:

- ◆ [“Unified View for Devices” on page 916](#)
- ◆ [“Unified View for Segments” on page 917](#)

Unified View for Devices

You can view All, Manageable, or Unmanageable devices in this view. For a corresponding device type, a device is said to be manageable if the list of MIBs implemented by the device satisfies the Manageability_definition property in the unifiedview.ini file in the novell zenworks\mms\mwserver\bin directory. The Manageability_definition property can be updated with a valid boolean expression of MIB names.

Following are the device types that you can filter:

- ◆ All (all types of devices)
- ◆ NetWare Servers
- ◆ NCP Print Servers
- ◆ TCP Services
- ◆ Printers
- ◆ IP Routers
- ◆ Switches/Bridges
- ◆ IPX Routers
- ◆ Windows Servers

To filter the devices:

- 1** In the atlas, select *View > Unified View for Devices*.
- 2** From the first drop-down list, select *All* to list all the devices.
or
Select *Manageable* to list the manageable devices.
or
Select *Unmanageable* to list the unmanageable devices.
- 3** From the second drop-down list, select a device type.
- 4** Click *Show*.

The Unified view will display the list of the devices. The tabular column in the Unified view contains the following information:

- ◆ The icons associated with the devices.
- ◆ The MIBs implemented by the device. If the device does not implement any MIBs the column will specify “No MIBs implemented” for that device.
- ◆ The maximum severity of the alarms against the devices. To view the legend for the alarm, select the alarm legend button on the toolbar.

Unified View for Segments

You can view All, Manageable, or Unmanageable segments in this view. For a corresponding segment type, a segment is said to be manageable if the list of MIBs implemented by at least one device in that segment satisfies the `Manageability_definition` property in the `unifiedview.ini` file. The `Manageability_definition` property can be updated with a valid boolean expression of MIB names.

The following are the segment types you can set filter for:

- ◆ All (all types of segments)
- ◆ Ethernet
- ◆ Frame Relay
- ◆ IPX Compatibility Mode

- ◆ Token Ring
- ◆ X.25
- ◆ PPP
- ◆ ATM
- ◆ FDDI

To filter the segments:

- 1** At the Atlas level, select *View > Unified View for Segments*.
- 2** From the first drop-down list, select *All* to list all the segments.
or
Select *Manageable* to list the manageable segments.
or
Select *Unmanageable* to list the unmanageable segments.
- 3** From the second drop-down list, select a segment type.
- 4** Click *Show*.

The Unified view will display the list of the segments. The tabular column in the Unified view contains the following information.

- ◆ The icons associated with the segments.
- ◆ The name of the segment.
- ◆ The maximum severity of the alarms against the segments. To view the legend for the alarm, select the alarm legend button on the toolbar.

Understanding Alarm Management

24

The Novell® ZENworks® Server Management Alarm Management System alerts you to important events like the SNMP traps, threshold alarms, and ping and connectivity testing faults occurring on your network. This lets you proactively resolve network problems and receive updates on events occurring on your network.

Alarm icons are anchored to objects displayed in Novell ConsoleOne®. The icons change color to depict the level of severity, notifying you of potential problems. The events are reported in the Active Alarm view, and each event is categorized and displayed with a corresponding alarm icon.

The Alarm Management System processes any device on the network that supports SNMP-standard trap notification. For example, for all Novell NetWare® servers on which the Management Agent for NetWare is installed, notifications of server breakdowns, overloads, and configuration changes are sent to the management server for processing and then made available for viewing at a Novell ZENworks Server Management Novell ConsoleOne.

You can enable and disable alarms and set alarm thresholds on baseline statistics for segments and servers (for example, segment alarms for utilization and the total number of packets per second), so that an alarm is generated when the threshold for a statistic is reached. You can also set actions to be performed when an alarm occurs. The actions assigned to an alarm are specified in the alarm rule.

This section contains the following topics:

- ◆ [Section 24.1, “Understanding the Alarm Management System,” on page 919](#)
- ◆ [Section 24.2, “Managing the Alarm Management System,” on page 924](#)
- ◆ [Section 24.3, “Managing the Rule-Based Alarm Management System,” on page 935](#)
- ◆ [Section 24.4, “Maintaining the Alarm Management System,” on page 952](#)
- ◆ [Section 24.5, “Troubleshooting the Alarm Management System,” on page 952](#)

24.1 Understanding the Alarm Management System

The Alarm Management System alerts you to network conditions and events. It provides you with tools and back-end services to use, distribute, and manage this information. The Alarm Management System component is also fully integrated with other Novell ZENworks Server Management components. It provides access control through the Role-Based Services component and provides report generation through the reporting functions. The Alarm Management System provides a centralized location for processing and viewing the events and alarms generated by devices and systems throughout your network.

You can use ConsoleOne to view tabular lists of statistical data for active and historical alarms received by the Alarm Management System. This makes it easy to handle alarms and track network events and recurring alarm conditions.

In addition, real-time notification of alarms occurring on your network is provided by the following:

- ◆ Severity level, as displayed by the changing color of the alarm indicators
- ◆ Audible notification
- ◆ Status bar ticker-tape messages

You can also assign an action to an alarm, such as automatically launching a program when an alarm is received, or sending an e-mail message to notify remote users of events.

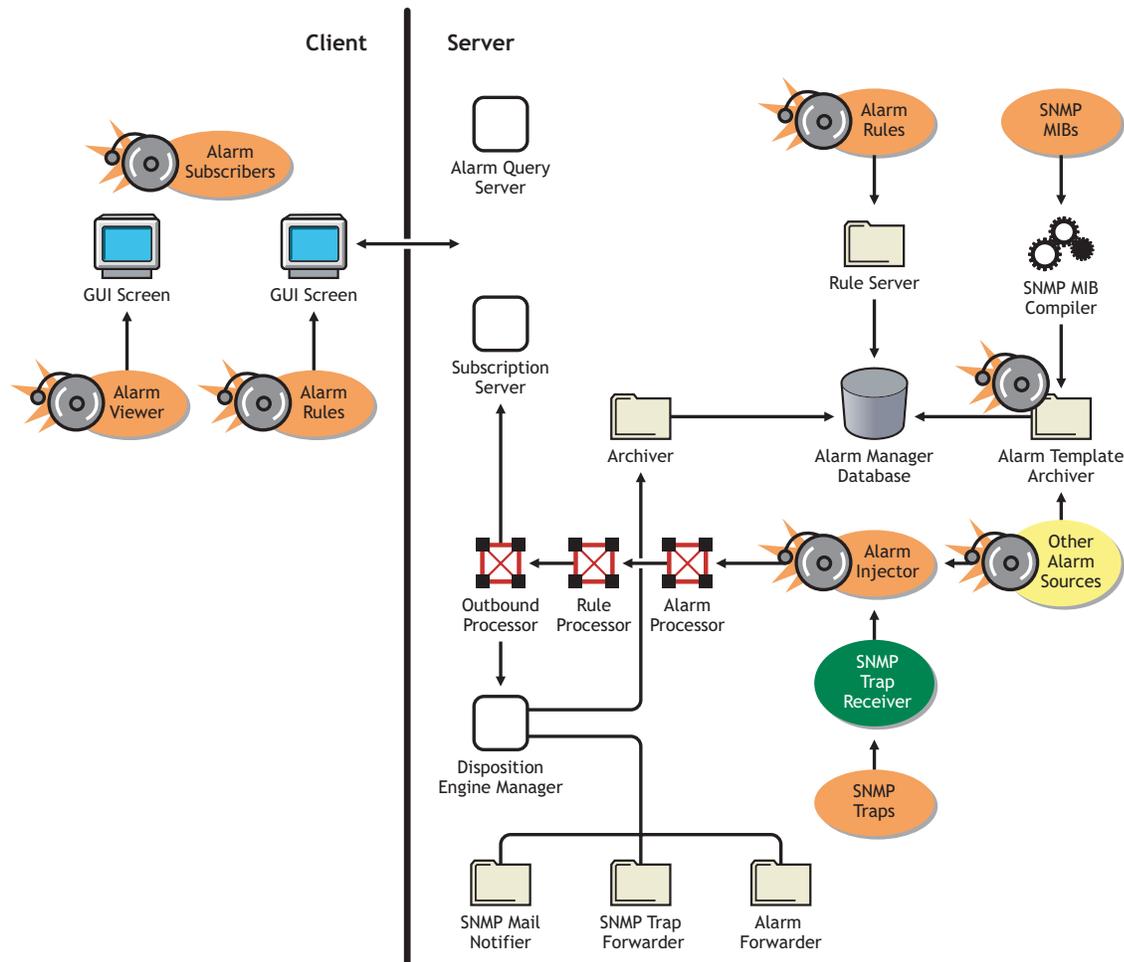
- ◆ [Section 24.1.1, “Alarm Management System Components,” on page 920](#)

24.1.1 Alarm Management System Components

The Alarm Management System consists of multiple components for processing, storing, and viewing alarms. All alarms received by the Alarm Management System are processed and sent to applications that subscribe to them. ConsoleOne, by default, subscribes to the Alarm Management System and receives updates when an alarm is processed. Hierarchical Status Notification also subscribes to the Alarm Management System and changes the color of the atlas map icon accordingly.

[Figure 24-1](#) illustrates the Alarm Management System components:

Figure 24-1 Alarm Management System components on the client and on the server



The main components that make up the Alarm Management System are as follows:

- ◆ “SNMP Trap Receiver” on page 921
- ◆ “Alarm Server” on page 922
- ◆ “Alarm Processors” on page 922
- ◆ “Rule Processor” on page 922
- ◆ “Outbound Processor” on page 922
- ◆ “Disposition Engine Manager” on page 922
- ◆ “Alarm Manager Database” on page 922
- ◆ “Archivers” on page 923
- ◆ “Alarm Viewers” on page 923

SNMP Trap Receiver

The SNMP Trap Receiver receives traps from network management agents and converts them to alarms. Thereafter, it passes them to the Alarm Server.

Alarm Server

The Alarm Server receives alarms from the SNMP Trap Receiver and other applications. Then it passes them to the alarm processors.

Alarm Processors

The Alarm Processors include processes for receiving, processing, and dispatching alarms to various subscribers. The inbound processor applies alarm templates to incoming alarms. After inbound processing is completed, the alarm is sent to the LSM Hook Processor, which processes segment-related alarms. The LSM Hook Processor sends the alarm to the Rule Processor.

Rule Processor

The Rule Processors applies all the configured rules on the alarm it has received from the LSM Hook Processor. If the alarm satisfies any of the rules, corresponding disposition information is updated and the alarm is sent to the Outbound Processor. If none of the rules satisfies by the alarm, the alarm is dropped.

Outbound Processor

After receiving the alarm from the Rule Processor, Outbound Processor sends the alarm to the Disposition Engine Manager and all the subscribers of the alarm manager.

Disposition Engine Manager

The Disposition Engine Manager checks the Actions defined in the alarm it has received and sends the alarms to respective disposition engines such as the SNMP Trap Forwarder, Alarm Forwarder, Archiver, and SMTP Mail Notification.

Alarm Manager Database

The alarm manager database, a repository for alarm information, includes the following:

- ◆ [“Processed Alarms” on page 922](#)
- ◆ [“Alarm Templates” on page 922](#)
- ◆ [“Alarm Rules” on page 923](#)

Processed Alarms

The processed alarm data that is stored in the alarm manager database is supplied to ConsoleOne through the alarm query server. The alarm data is used for alarm and alarm summary presentation and reporting.

Alarm Templates

Alarm Templates are applied to each alarm received by the inbound processor. The alarm template is based on SNMP trap definitions in the MIB or other proprietary definitions for handling the Alarm Management System management and display criteria. When you compile the MIB, the trap definitions are used to create an alarm template that provides a method for presenting and managing alarm data. Proprietary alarm templates are based on proprietary definitions.

For example, when a user tries to log in to a server with an incorrect password, an alarm is generated and forwarded to the management server. The management server processes the alarm by identifying the trap object identifier (OID) and assigns the associated alarm template.

A default template is assigned to an SNMP trap sent by a device that does not have a recognizable OID and is categorized as unknown. In order for a trap OID to be recognized by Alarm Management System, you need to compile the MIB of the device into the MIB Pool on the management server.

Alarm Rules

Alarm Rules govern the handling characteristics of SNMP traps or proprietary alarms. Each Alarm Rule contains a set of Conditions and Actions. For example, Source address, Alarms, Severity, State, and Time Interval are Conditions and Sending SMTP Mail Notification, Trap Forwarding, Archiving, Launching Applications, and Automatic Assignment to User are Actions. An alarm can only satisfy a Rule when it complies to all Conditions and Actions specified in the Rule. When an alarm satisfies a rule, the Actions defined in the Rule perform the specified operations on the alarm.

Archivers

The following three archivers add data to the alarm manager database:

- ◆ [“Alarm Archiver” on page 923](#)
- ◆ [“Rule Server” on page 923](#)
- ◆ [“Template Archiver” on page 923](#)

Alarm Archiver

The alarm archiver stores alarm statistics and data in the alarm database. By default, all alarms are archived. If you do not want an alarm to be archived, you can disable the default rule. See [“Archiving Alarm Statistics” on page 945](#) for more information.

Rule Server

The Rule Server receives the alarm rule from the Alarm Rule Console and saves it in the alarm manager database.

Template Archiver

The template archiver receives alarm templates from a MIB compiler and saves them in the alarm manager database.

Alarm Viewers

ConsoleOne displays three views of alarm data: the Active Alarm view, the Historical Alarm view, and the Alarm Summary view.

The Active Alarm view displays statistics in ConsoleOne for events occurring on your network. Alarms displayed in the Active Alarm view can either be owned by you or assigned to a group. The tasks that you can perform on an alarm from this view depend on the access rights allowed through the Role-Based Services. The Active Alarm view appends incoming alarms to the list, providing you with the most recent alarms. After an alarm is handled, it is removed from the Active Alarm list.

The Alarm History view displays information about assignments and ownership of alarms. You can track alarms received by the Alarm Management System and verify their handling status from this view.

The Alarm Summary view is a graphical representation of all the alarms that you have received.

24.2 Managing the Alarm Management System

ConsoleOne provides a central location for monitoring, managing, and controlling critical events on your network. You can configure the Alarm Management System to alert you to errors on critical systems and events to assist you in maintaining your network. This section contains the following information:

- ◆ [Section 24.2.1, “Recognizing Alarm Indicators,” on page 924](#)
- ◆ [Section 24.2.2, “Viewing Alarms,” on page 925](#)
- ◆ [Section 24.2.3, “Enabling and Disabling Alarms,” on page 929](#)
- ◆ [Section 24.2.4, “Resolving Alarms,” on page 929](#)
- ◆ [Section 24.2.5, “Deleting Alarms,” on page 931](#)
- ◆ [Section 24.2.6, “Performing Actions on Alarm Templates,” on page 933](#)

24.2.1 Recognizing Alarm Indicators

You can monitor the network for alarm-triggering events by observing nodes on topology maps or Atlas views, [Active Alarm](#), and [Alarm History](#) views, and in the server/node summary. [Table 24-1](#) lists the alarm indicators and the type of alarm they are associated with:

Table 24-1 Alarm indicators and its associated Alarm type

| Alarm Indicator | Applies To |
|---|---|
| Alarm icons anchored to the affected object | Severe, major, and minor alarms are displayed in the Atlas and Console views and the left pane of ConsoleOne. An alarm icon remains anchored to a segment or device object until you handle all alarms outstanding against that object. Alarm icons differ based on the severity level of the alarm. See “Interpreting Alarms” on page 926 for details on alarm severity and the associated icons. If a segment or device has multiple alarms logged against it, the alarm icon always depicts the highest level of severity. |
| Ticker-tape message on the status bar | By default, the Alarm Management System automatically displays alarm messages on the status bar. You can enable or disable this ticker-tape display for each severity level. For information on configuring this option, see “Displaying a Ticker-Tape Message” on page 945 . |
| Audible beep | The Alarm Management System can be configured to produce an audible beep on ConsoleOne when an alarm occurs. By default, this option is disabled. You can configure each individual severity level to enable the audible notification. For information on setting this option, see “Beep On Console” on page 946 . |

24.2.2 Viewing Alarms

You can access active and historical alarm data from any ConsoleOne location. As an administrator, you can define access restrictions to alarm data and management functions through Role-Based Services to further define the data presented based on the roles in your organization.

You can modify the presentation of the alarm data displayed in the Active Alarms and Alarm History view by filtering the displayed data, changing the column layout, and changing the sorting order. All options for changing the presentation are under the View menu in ConsoleOne.

The following sections describe the different ways you can view and use alarms:

- ♦ [“Viewing Active Alarms” on page 925](#)
- ♦ [“Viewing Historical Alarms” on page 925](#)
- ♦ [“Viewing the Alarm Summary” on page 926](#)
- ♦ [“Interpreting Alarms” on page 926](#)
- ♦ [“Sorting Alarms” on page 927](#)
- ♦ [“Filtering Alarms” on page 928](#)

Viewing Active Alarms

The ConsoleOne Active Alarm view displays alarm statistics for all current alarms received from segments or devices, per management domain. The Summary view shows a list of all active alarms for that server or node.

The Active Alarms view and Server Summary view display a table of detailed information about active alarms. These views are updated whenever a new alarm occurs and is archived on your network. New alarms are appended to the list.

To display the Active Alarm view:

- 1 In ConsoleOne, select the Novell ZENworks Server Management site object.
- 2 Click *View > Active Alarms*.

The Active Alarm view is displayed. You can perform the following activities from this view:

- ♦ [“Assigning Alarms” on page 930](#)
- ♦ [“Owning Alarms” on page 930](#)
- ♦ [“Handling Alarms” on page 931](#)
- ♦ [“Adding Notes to Alarms” on page 931](#)

Viewing Historical Alarms

The Alarm History view displays information about all archived alarms, including the handling status of each alarm. You can access the Alarm History view only if you have been granted access through the Role-Based Services.

To display the Alarm History view:

- 1 In ConsoleOne, select the Novell ZENworks Server Management site object.
- 2 Click *View > Alarm History*.

The Active Alarm view is displayed. You can perform the following alarm handling activities from this view:

- ◆ “Assigning Alarms” on page 930
- ◆ “Owning Alarms” on page 930
- ◆ “Deleting Alarms” on page 931
- ◆ “Adding Notes to Alarms” on page 931

Viewing the Alarm Summary

The Alarm Summary is a graphical representation of the summary of alarms you have received. The view is divided into three panels: a pie chart panel, a bar graph panel, and a trend panel. You can choose to view the information in these panels for a given period of time. The time duration is for the hour, for the day, for the week, and for the month.

- ◆ The pie chart panel includes alarm distribution based on severity, category, owner, and alarm state
- ◆ The bar graph panel includes the Top N Alarm types, Top N Source Address and Top N Affected Node. The value of N is configurable.
- ◆ The trend displays the rate at which the alarms are received.

You can customize the pie chart and the bar graph representations to reflect the customized data.

To display the Alarm Summary view:

- 1** In ConsoleOne, select the Novell ZENworks Server Management site object.
- 2** Click *View > Alarm Summary*.

To customize the pie chart and the bar graph representation:

- 2a** In the Alarm Summary view, click the *Customize* button to display the Customize Summary View dialog box.

By default, all the options in this dialog box are selected. You can select to display only the options you want.

Interpreting Alarms

The Active Alarm and Alarm History views display lists of alarms that have been archived in the alarm manager database. The alarms are displayed as a tabular list. [Table 24-2](#) describes the data types and contents:

Table 24-2 Alarm Data Types

| Data Type (Column) | Contents |
|--------------------|---|
| Severity | Alarm icon that indicates the severity level attributed to the trap. The color of the alarm icon indicates the level of alarm severity, as follows:  Red = Severe  Magenta = Major  Yellow = Minor  Blue = Informational  White = Unknown |
| From | Network address of the device that sent the alarm to the Alarm Management System. |
| Summary | Summary of the event, often including the name or address of the object affected by the alarm. |
| Owner | Person or group responsible for handling the alarm. The default owner is SYSTEM. |
| Received Time | Date and time when the Alarm Management System received the alarm. |
| Type | Generic description of the alarm. For example, volume out of disk space. |
| Category | Category identified in the MIB associated with the trap-type object. |

You can filter the data displayed in the alarm views based on criteria from statistics displayed in each view; see [“Filtering Alarms” on page 928](#) for details. After selecting one or more alarm entries in an alarm view, you can perform operations by right-clicking them.

Sorting Alarms

You can modify the order in which the alarms are displayed on the Active Alarm or Alarm History views by sorting the alarms. By default, the alarms are sorted in ascending order by received time.

To edit the sort settings:

- 1 Click *View > Settings > Sort*.
- 2 Select the criteria by which you want the alarms sorted. You can sort by
 - ◆ *Type*
 - ◆ *Severity*
 - ◆ *Category*
 - ◆ *Received time*
 - ◆ *Summary*
 - ◆ *Owner*
 - ◆ *Affected Object*
- 3 Indicate whether you want the alarms sorted in ascending (oldest first) or descending (the most recent alarms first) order by selecting the appropriate option in the *Sort Order* box.

- 4 Click *OK*.

Filtering Alarms

You can display the alarms in a tabular view based on filter conditions. The filter applies only to the current management session and clears when you exit ConsoleOne.

You set up a filter by selecting criteria from four drop-down lists. You can either set up simple filters that require only one line, or complex filters composed of multiple lines or groups of lines. If you set up a filter using more than one line, you must also specify the logical relationship between the line and/or group of lines.

To set up a filter:

- 1 Go to the view you want to filter.
- 2 Click *View > Settings > Filter*.

The Alarm Filter dialog box is displayed.

- 3 Select the column by which you want the Alarm Management System to filter alarms from the first drop-down list. You can filter alarms using the following columns:
 - ◆ **Severity**: Filters the alarms based on the alarm severity. Alarm severity is assigned to an alarm type.
 - ◆ **Generator Type**: Filters alarms based on the type of agent or system generating the alarms.
 - ◆ **Category**: Filters alarms based on the category of the alarm. Alarm categories are based on the MIB that defines the trap-type objects.
 - ◆ **Type**: Filters alarms based on the alarm type. The alarm type is set by the SNMP trap-type defined in the MIB or the proprietary alarm definition.
 - ◆ **Source Address**: Filters alarms based on the source addresses.
 - ◆ **Affected Object**: Filters alarms based on the affected objects.
 - ◆ **Alarm Owner**: Filters alarms based on the owner of the alarm.
 - ◆ **Alarm Summary**: Filters alarms based on the alarm summary.
 - ◆ **Rule**: Filters alarms based on the rules.
- 4 Select an operator from the second drop-down list.

The operator defines how to constrain the column you have selected to a value. For example, you can specify that the selected category must be equal to, not equal to, greater than, less than, greater than or equal to, less than or equal to, contain, or start with the value you select in the third drop-down list in order for an alarm to be displayed. Keep in mind that the list of available operators depends on the column you've selected.
- 5 Select a value from the third drop-down list.
- 6 Specify how this filter statement relates to other statements you plan to define by selecting a value from the fourth drop-down list.
 - ◆ If this is the only filter statement or if it is the last statement in a group, select *End*.
 - ◆ If you want to add a line below the current filter statement, select *New Row*. A new line is added. You must define the logical relationship between the previous line and the new line. The alarms are displayed based on the logical condition you have specified. Select

And to satisfy both the filter conditions. Select *Or* to satisfy any one of the filter conditions for the alarm to be displayed.

- ◆ If you want to add one or more lines that are unrelated to the preceding lines, select *New Group*. A new line is added. An additional drop-down list separates the new line from the preceding lines. Select a value from this drop-down list to indicate the relations between the filter statements. Select *And* if you want both the filter statements to be satisfied. Select *Or* if you want only one of the filter statements in one of the groups to be satisfied. Select *End* from the fourth drop-down list when you add a new group.

7 Click *OK* if you have defined filters.

The alarm list is updated to display only those alarms that meet the filter criteria you defined.

24.2.3 Enabling and Disabling Alarms

Novell ZENworks Server Management provides default threshold values for managed NetWare and Windows* servers and network segments hosting the Traffic Analysis Agents for a station connected to a segment. An alarm is generated if the values exceed the threshold values. The server threshold alarms are enabled by default, but the segment threshold alarms are not.

IMPORTANT: In order to modify the segment properties, you must have the Traffic Analysis Agents for NetWare or Windows hosted on a station, connected to the segment.

To enable or disable segment threshold alarms:

- 1 Right-click the segment object, then click *Properties*.
- 2 If it is not already displayed, select the *Segment Alarms* tab.
- 3 Select the alarm you want to enable or disable, then click *Edit*.
- 4 In the *Value* field, specify the threshold value after which an alarm should be generated.
- 5 In the *Sampling Interval* field, specify the time (in seconds) that the threshold value must exceed in order to generate an alarm.
- 6 Select the *Enable* check box.
- 7 Click *OK*.

24.2.4 Resolving Alarms

Alarms that occur on segments and devices on your network are added to the alarm manager database and are presented in the Active Alarms and Alarm History views. Entries in the alarm manager database remain in the database until the alarm is deleted. The database records the status of the alarm as it is first acknowledged, then assigned to a group or user, to the point where it is deleted from the database after the owner has resolved the problem.

Resolution operations for alarms are displayed when you right-click a single entry or multiple entries in an alarm view and click any of the following actions:

- ◆ [“Assigning Alarms” on page 930](#)
- ◆ [“Owning Alarms” on page 930](#)
- ◆ [“Handling Alarms” on page 931](#)
- ◆ [“Adding Notes to Alarms” on page 931](#)

- ◆ [“Jump to the Affected Node” on page 931](#)

You can also access the alarm action menu items from the *View* menu in Novell ConsoleOne.

The order in which you perform the handling, assigning, and owning of an alarm or multiple alarms depends on your organization. Keep in mind that after you handle an alarm, it is removed from the Active Alarms list and only appears in the Alarm History list. A suggested course for resolving an alarm is for you to first assign the alarm to a group or team member, then have someone from the group take ownership of the alarm. When the network problem or event has been resolved, the team member can remove it from the Active Alarms list and eventually delete it from the Alarm History. By following this process, you can track the alarm status from creation through resolution, until it is finally deleted from the Alarm History list.

Assigning Alarms

You can specify the group or user that is assigned to handle an alarm. This allows you to use any team assignments you already have within your organization. For example, you might have a group or team member assigned to handle all alarms relating to NetWare servers. You can assign one or more alarms to a group or user if you have been granted access to assign alarms through the Role-Based Services. You can use an alarm filter to help you determine groups based on certain filtering criteria. See [“Filtering Alarms” on page 928](#) for information on filtering options.

To assign an alarm:

- 1 Select the alarm you want to assign from the [Active Alarm](#) or [Alarm History](#) list.
- 2 Click *View > Assign*.
- 3 In the *Username* field, select the name of the person or group to which you want to assign the alarm.

For more information on users, see [“Adding a New User” on page 950](#)

The name you select does not correlate to users in Novell eDirectory and can represent the organizational structure you already have in place.

- 4 Click *OK*.

Owning Alarms

A user can take ownership of one or more alarms. If a user is a member of a group assigned to resolve a network problem, the team member can take ownership of the alarm and eventually delete the alarm to remove it from the alarm manager database.

To take ownership of an alarm:

- 1 Select the alarm from the [Active Alarm](#) or [Alarm History](#) view.
- 2 Click *View > Own*.

The value in the *Owner* field changes to the Novell eDirectory name you are logged in as.

You cannot customize this option; the user logged in to ConsoleOne always becomes the owner of the alarm when this action is used.

Handling Alarms

Alarms displayed in the Active Alarm view have not been handled by anyone. After the alarm is handled, it is removed from the *Active Alarm* list, and any alarm indicators shown in other views in ConsoleOne are removed.

See “[Recognizing Alarm Indicators](#)” on page 924 for information on different types of alarm indicators. After it is removed from the *Active Alarm* list, the alarm is still displayed in the [Alarm History](#) view until it is deleted by the owner.

To handle an alarm:

- 1 Select the alarm from the [Active Alarm](#) list.
- 2 Click *View > Handle*.

The alarm is removed from the Active Alarm list. You can still display information about the alarm by switching to the [Alarm History](#) view.

Adding Notes to Alarms

You can add a note to any of the alarms displayed in the Active Alarm view or Alarm History view. The note can contain any relevant useful information about the alarm.

- 1 Select the alarm from the [Active Alarm](#) or [Alarm History](#).
- 2 Click *View > Note*.
- 3 In the Note dialog box, create a note for the alarm.
- 4 Click *OK*.

The alarm icon now has a note icon associated with it, indicating that a note has been added to the alarm.

If you want to delete the note from the alarm, repeat step 2. Delete the note that you created in the Note dialog box, then click *Apply*. The note is deleted for the alarm, and the note icon is not displayed.

Jump to the Affected Node

You can jump to the affected node where the alarm has been triggered and perform the necessary action to rectify the problem.

- 1 Select the alarm from the [Active Alarm](#) or [Alarm History](#).
- 2 Click *View > Jump to Affected Node*.

The Console view is displayed and the node on which the alarm has triggered is highlighted.

24.2.5 Deleting Alarms

Alarms displayed in the Alarm History view can be deleted from the alarm list after problem resolution. You can delete one or more alarm entries to remove the alarm from the list. To delete an alarm, you must have access to view the alarm history and to delete alarms through the Role-Based Services.

There are two ways to delete alarms:

- ◆ You can delete alarms manually from the Alarm History view. See [“Deleting Alarms from ConsoleOne” on page 932](#)
- ◆ You can delete alarms automatically using the Alarm Management System purge utility. See [“Deleting Alarms Using the Purge Utility” on page 932](#).

IMPORTANT: The alarm manager database, located on the management server, records the status of every alarm instance received by the Alarm Management System. You must be diligent in deleting alarms after a problem is resolved in order to keep the database from taking up excessive disk space. Currently, the alarm manager database uses the Alarm purge utility (on by default) to automatically delete entries after a period of time or based on the size of the database.

Deleting Alarms from ConsoleOne

You can manually delete alarms through ConsoleOne.

- 1 Select the alarms you want to delete from the **Alarm History** list.
- 2 Click *View > Delete*.

The alarms are removed from the Alarm History view.

Deleting Alarms Using the Purge Utility

You can delete alarms automatically using the Alarm Management System purge utility. Before you can use this utility, you must set up the utility's configuration file, `ampurge.properties`, which is located in the properties directory on the server and volume where you installed the alarm manager database. Then you can schedule the utility to run automatically at a specified time of day. The following sections describe how to set up and use the Alarm Management System purge utility:

- ◆ [“Setting Up the Purge Utility Configuration File” on page 932](#)
- ◆ [“Setting Up the Purge Utility to Run Automatically” on page 933](#)

Setting Up the Purge Utility Configuration File

The Alarm Management System purge utility configuration file, `ampurge.properties`, defines the criteria for selecting the alarms to be purged as well as the time of day the process should run. This file is located in the properties directory on the server and volume where you installed the alarm manager database.

Before you can run the purge utility, you must set up the configuration file as follows:

- 1 Open the `ampurge.properties` file with a text editor.
- 2 Set the criteria for purging alarms by editing the values of the following lines in the file:
 - ◆ `SeverityInformationalPurgeWait`: The number of days before informational alarms will be purged.
 - ◆ `SeverityMinorPurgeWait`: The number of days before minor alarms will be purged.
 - ◆ `SeverityMajorPurgeWait`: The number of days before major alarms will be purged.

- ♦ `SeverityCriticalPurgeWait`: The number of days before severe alarms will be purged.
- ♦ `SeverityUnknownPurgeWait`: The number of days before unknown alarms will be purged.

By default, alarms of all severity levels are purged after seven days.

- 3 Save the configuration file.

Setting Up the Purge Utility to Run Automatically

You can schedule the purge utility to run daily to ensure that the alarm manager database does not consume excessive disk space. Before you can set up the utility to run automatically, you must make sure to set up the file with your preferences for deleting alarms of various severities. See [“Setting Up the Purge Utility Configuration File” on page 932](#).

To set up the utility to run automatically:

- 1 Open the `ampurge.properties` file with a text editor.
- 2 Set the time of day you want the utility to run by editing the `PurgeStartTime` entry.
Valid values are 0 to 23, where 0 is midnight and 23 is 11:00 p.m. Keep in mind that the purge utility is memory intensive and can occupy the server for several minutes. Therefore, you should set the utility to run during off-peak hours.
- 3 Save and close the file.
- 4 Open the `alarmmanager.properties` file and verify that the following line exists:
`AlarmPurgeService=yes`
If the line does not exist, add it to the end of the file.
- 5 Save and close the file.
- 6 Restart the server.

24.2.6 Performing Actions on Alarm Templates

By editing the alarm disposition associated with each alarm template, you can configure an alarm to automatically perform an action when an alarm occurs. Alarm dispositions are created for each alarm template in the Alarm Manager database and default settings are assigned to them. You can edit the alarm dispositions to enable the following actions:

- ♦ [“Sorting Alarm Templates” on page 933](#)
- ♦ [“Modifying the Severity and State of the Alarm” on page 934](#)
- ♦ [“Deleting Alarm Templates from Novell ConsoleOne” on page 935](#)
- ♦ [“Printing the Alarm Disposition” on page 935](#)
- ♦ [“Copying Alarm Templates to Microsoft Excel” on page 935](#)

Sorting Alarm Templates

The Alarm Management System system enables you to sort the alarm templates based on different conditions. This option is enabled by default. You can sort the templates based on Severity, Generator Type, Category, or Type. By default, the sorting is done based on the Type. You can also sort the templates based on a single field by selecting the field from the drop-down list under the

Sort Items By option, or you can sort the templates based on different combinations of fields by using the *Then By* options.

- 1 In ConsoleOne, right-click the Novell ZENworks Server Management site object, then click *Properties*.
- 2 Click the *Alarm Disposition* tab.
- 3 In *Templates* tab, click the *Sort* button to display the Template Sorting dialog box.
- 4 Select fields from *Sort Items By* drop-down list.
- 5 Select fields from *Then By* drop-down list.
- 6 Click *OK*.

The templates are sorted based on the field selected in the *Sort Items By* and *Then By* options. For example, if you want to sort the templates based on severity, category, generator type, and Type, first select *Severity* in the *Sort Items By* list, then select *Category*, *Generator Type* and *Type* in the three *Then By* drop-down lists. The templates are sorted first based on severity, then on the category, then by the generator type, followed by the type.

Modifying the Severity and State of the Alarm

The alarm disposition includes other configuration settings that include modifying the severity or the state of the alarm. You can modify the severity or the state of the alarm, or both. The incoming alarms will display the modified severity and the state.

To change the severity and/or state of the alarm:

- 1 In ConsoleOne, right-click the Novell ZENworks Server Management site object, then click *Properties*.
- 2 Click the *Alarm Disposition* tab.
- 3 In the *Templates* page, select the alarm that you want to edit from the *Alarm Templates* list, then click *Modify* to display the Modify Template dialog box.
- 4 To modify the severity of the alarm, select the severity from the drop-down list. You can change the severity of the alarms to one of the following:

- Informational
- Minor
- Major
- Critical
- Default

- 5 To modify the state of the alarm, select the state from the drop-down list. You can change the the alarm state to one of the following:

- Operational
- Degraded
- Non-operational
- Default

- 6 Click *OK*.

Deleting Alarm Templates from Novell ConsoleOne

You can now delete the alarm templates from the Alarm Templates list through ConsoleOne. The alarms corresponding to the deleted templates are not processed by the Alarm Manager.

- 1 In ConsoleOne, right-click the Novell ZENworks Server Management site object, then click *Properties*.
- 2 Click the *Alarm Disposition* tab.
- 3 In *Templates* tab, select the alarm template that you want to delete from the *Alarm Templates* list, then click *Delete*.
- 4 Click *Yes* in Delete Template dialog box.

Printing the Alarm Disposition

- 1 In ConsoleOne, right-click the Novell ZENworks Server Management site object then, click *Properties*.
- 2 Click the *Alarm Disposition* tab.
- 3 On the *Templates* page, select the alarm template from the *Alarm Templates* list, then click *Print*. You can also select multiple alarm templates by pressing Ctrl + clicking them.
- 4 Select the alarms in the Print Alarms dialog box, then click *Print*.

Copying Alarm Templates to Microsoft Excel

You can now copy the alarm templates to a Microsoft* Excel spreadsheet.

- 1 Select the alarm templates you want to copy.
- 2 Press Ctrl+C.
- 3 In the Microsoft Excel spreadsheet, press Ctrl+V.

24.3 Managing the Rule-Based Alarm Management System

The Rule-Based Alarm Management System allows you to configure a rule on the site server and helps you in processing, storing, and monitoring alarms. It provides you with tools and back-end services to use, and manage the rule information. The Rule-Based Alarm Management System component is also fully integrated with other Novell ZENworks Server Management components. It provides access control through the Role-Based Services component. The Rule-Based Alarm Management System provides a centralized location for processing and viewing the rules generated by devices and systems throughout your network. In addition, you can also set specific conditions and actions to be performed when an alarm occurs.

A rule consists of two parts, namely set of conditions and set of actions. When Alarm Manager receives an alarm, it matches various parameters of the Alarm with the set of conditions in a rule. If all the conditions specified in a rule are satisfied, then the set of actions specified in the rule is performed. The rules are processed in the order of precedence and the topmost rule has the highest precedence.

When more than one rule is specified, the incoming alarm is matched with the rules beginning with the topmost enabled rule in the list. If the incoming alarm matches with any rule, then only

corresponding actions are carried out, and there will be no further processing of remaining rules. If the incoming alarm fails to satisfy any of the enabled rules in the list, then the Alarm is discarded.

ConsoleOne provides a mechanism to organize the alarms based on the Alarm Disposition Rules.

For more information on the Alarm Disposition Rule, see the following topics:

- ◆ [Section 24.3.1, “Understanding the Properties,” on page 936](#)
- ◆ [Section 24.3.2, “Understanding the Conditions,” on page 936](#)
- ◆ [Section 24.3.3, “Understanding the Actions,” on page 940](#)
- ◆ [Section 24.3.4, “Performing Actions on Rules,” on page 947](#)

24.3.1 Understanding the Properties

The Properties page available from the of Rule Configuration page contains the name of the rule and its corresponding description. The description that you specify here is displayed in the Rule Description on the Available Rules page.

The Properties page also displays the date and time when the rule was created and when it was last modified. This information is system-generated.

The title bar of the Rule Configuration page is appended with name of the rule you have provided in the Rule Name text box.

24.3.2 Understanding the Conditions

The Conditions page available from the Rule Configuration page consists of multiple conditions that you can define in order to process an incoming alarms. To configure a rule, you must define at least one Condition. You receive only those alarms that meet the conditions you have defined in the Conditions page.

The Conditions page includes the following conditions:

- ◆ [“Source Addresses” on page 936](#)
- ◆ [“Severity, State, and Specific Alarms” on page 937](#)
- ◆ [“Time Intervals” on page 939](#)

Source Addresses

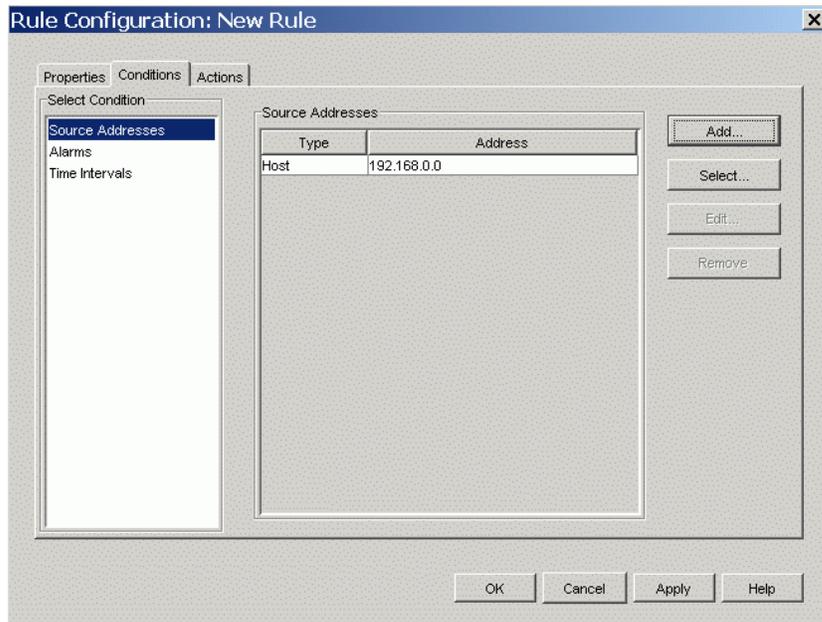
Source Address conditions enable you to process on incoming alarms based on their source addresses. You can define source addresses as either the host address or a range of IP addresses. You can also select one or more nodes or containers from the Atlas, which will be configured as a source address. The Source Addresses list is updated to display the host address or range of IP addresses or nodes you have added.

If you configure a condition to process an incoming alarm based on their source addresses, you receive only those alarms that are coming from the specified addresses.

To configure a rule with Source Address as a Condition:

- 1 Right-click the ZENworks Server Management site object in the left frame of ConsoleOne, then click *Properties*.

- 2 By default, Alarm Disposition page is selected. Click *Rule*, and then click *New* or *Edit*.
- 3 Click *Conditions*.
- 4 By default, the Source Address page is selected.



- 5 Click *Add*.
- 6 In Add Source Address dialog box, specify the Host Address, DNS name, or range of IP addresses, then click *OK*.
You can select one or more nodes or containers from the Atlas by clicking on *Select*.
- 7 Click *Apply* or *OK*.

As an administrator, you can also edit or remove the existing source address details.

NOTE: If you do not specify any source address, the system accepts all alarms, irrespective of the source address it comes from.

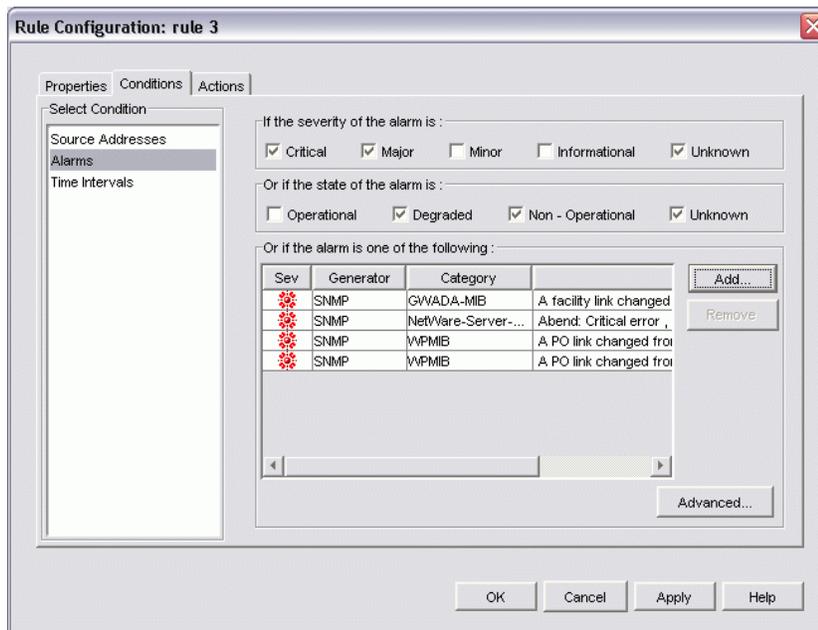
Severity, State, and Specific Alarms

The Conditions page of Rule Configuration page includes other configurations that filters alarms based on their severity, state, and ID. You can select one or more levels of severity, including Critical, Minor, Major, Informational, and Unknown. You can also select one or more state, including Operational, Degraded and Non-Operational. For example, if you select the severity as Critical and the state as Degraded, then you get alarms only if they are of Critical severity with the state as Degraded.

To configure a rule with specific alarms as a Condition:

- 1 Right-click the ZENworks Server Management site object in the left frame of ConsoleOne, then click *Properties*.
- 2 By default, Alarm Disposition page is selected. Click *Rule*, and then click *New* or *Edit*.

3 Click *Conditions*, then click *Alarms*.



4 Select one or more severity and state.

5 If you want to configure alarm as part of Conditions, click *Add*.

6 Select the alarms you want as part of Conditions.

You can also sort the alarms. While sorting the alarms, you can select one or more alarms by pressing **Ctrl +** clicking the alarms. For more details on sorting the alarms, refer to [“Sorting Alarm Templates” on page 933](#)

7 Click *Apply* or *OK*.

As an administrator, you can also remove the existing alarms from the list.

NOTE: You can also configure the alarms based on varbinds. This is an optional feature and is provided as an advanced configuration criteria. For more details on advanced alarm configuration, refer to [“Advanced Alarm Configuration” on page 938](#) below.

Advanced Alarm Configuration

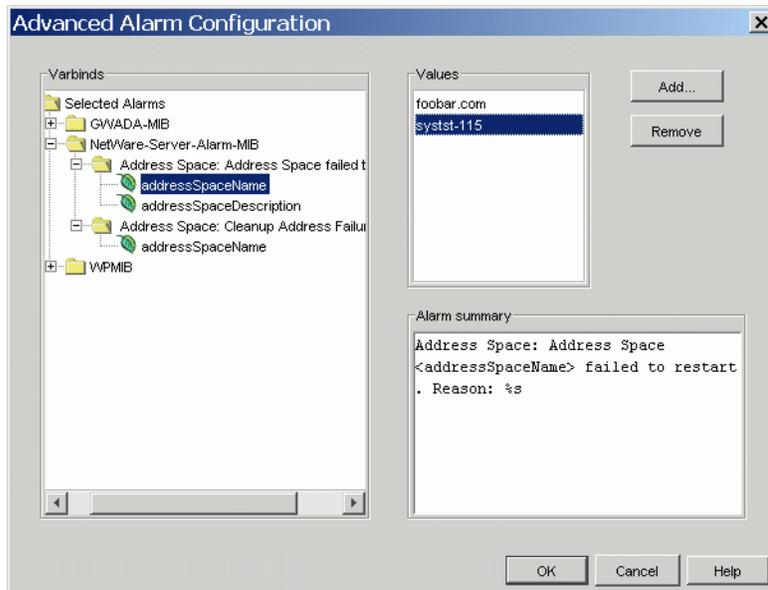
Advanced Alarm Configuration allows you to filter and select alarms based on their varbinds. Varbinds are the predefined variables of a trap. When a trap is generated, corresponding values are filled in the predefined variables and sent along with trap.

You can also add a value to the varbinds. To add a value, you must expand the tree until you see the varbinds of a selected alarm and select the required varbind.

To configure a rule with specific varbinds as a Condition:

- 1 Right-click the ZENworks Server Management site object in the left frame of ConsoleOne, then click *Properties*.
- 2 By default, Alarm Disposition page is selected. Click *Rule*, and then click *New* or *Edit*.

- 3 Click *Conditions*, then click *Alarms*.
- 4 Click *Add* and select an alarm that you want to set as part of *Conditions*.
- 5 Click *Advanced Alarm Configuration*.



- 6 In *Advanced Alarm Configuration* dialog box, expand the tree until you see the varbinds of a selected alarm.
- 7 Select a varbind to which you want to add a value.
- 8 Click *Add* and specify a value in *Add Value* dialog box, then click *OK*.
The *Advanced Alarm Configuration* page also displays summary of the alarm you have selected. The Summary contains name of the alarm and the varbinds included in it.
- 9 Click *OK* in the *Advanced Alarm Configuration* dialog box.
- 10 Click *Apply* or *OK*.

Varbinds values displayed in the list can also be deleted.

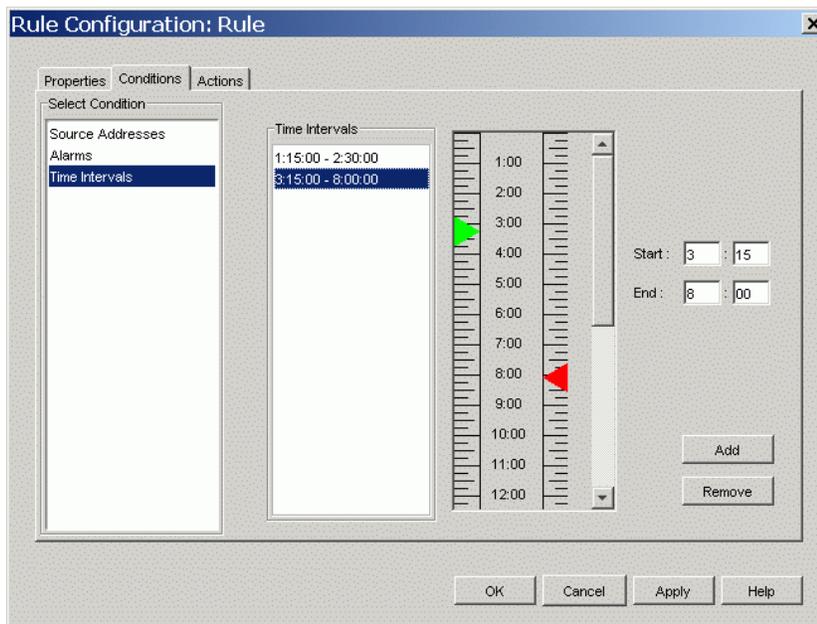
Time Intervals

You can also specify a time interval as a configuration criteria. If you do so, you will get an alarm only if the alarm arrives within the stipulated time interval.

To configure a rule with specific time interval as a Condition:

- 1 Right-click the ZENworks Server Management site object in the left frame of *ConsoleOne*, then click *Properties*.
- 2 By default, *Alarm Disposition* page is selected. Click *Rule*, and then click *New* or *Edit*.

3 Click *Conditions*, then click *Time Intervals*.



4 Move Green ruler to your desired time, which indicates starting time of an interval. Start time is set according to the ruler you have set.

You can also specify the time in *Start time* field. The ruler will set accordingly.

5 Move red ruler to your desired time, which indicates ending time of an interval. End time is set according to the ruler you have set.

You can also specify the time in *End time* field. The ruler will set accordingly

6 Click *Add*.

7 Click *Apply* or *OK*.

As an administrator, you can also remove the existing time interval. If you do not specify any time intervals, the system accepts all alarms, irrespective of the time when they come.

IMPORTANT: The time interval is always displayed in multiples of 15. For example, if you select the Start time as 1:10, the Alarm Management System takes the time as 1:15, and if you select the Start time as 1:30, the Alarm Management System takes the time as 1:30.

24.3.3 Understanding the Actions

Actions allows you to perform various actions when an alarm occurs. You can create a rule to automatically perform an action when an alarm occurs. To configure a rule, it is mandatory to define one of the actions.

The following sections describes the different ways you can configure Actions:

- ♦ “SMTP Mail Notification” on page 941
- ♦ “Launching an External Program” on page 942
- ♦ “SNMP Traps Forwarding” on page 944

- ♦ “Alarms Forwarding” on page 944
- ♦ “Miscellaneous” on page 945

SMTP Mail Notification

SMTP Mail Notification allows you to send SMTP messages to recipients who are specified to receive e-mail notification. You can also send SMTP mail notification to the SMTP server running on a port other than the default port.

The incoming and outgoing e-mail in the SMTP Host field is handled by the IP address or the port number of the SMTP host server.

To configure a rule with SMTP mail notification as an Action:

- 1 Right-click the ZENworks Server Management site object in the left frame of ConsoleOne, then click *Properties*.
- 2 By default, Alarm Disposition page is selected. Click *Rule*, and then click *New* or *Edit*.
- 3 Click *Actions*.
- 4 By default, the SMTP Mail Notification page is selected.
- 5 Select *Enable*.
- 6 Enter the IP address or the port number (optional) of the SMTP host server that handles incoming and outgoing e-mail in the *SMTP Host* field.

If you are specifying the port number, specify the port number after the IP address with a colon (:) preceding it. For example, if the IP address of the SMTP host server is 207.68.173.245, and the port number is 12345, specify as 207.68.173.245:12345.

You can click *Test* to verify if the SMTP server is running on the specified IP address or not.

- 7 Enter the e-mail address from where the mail notification is being sent.
- 8 Enter the e-mail addresses of the recipients in the *To* field.
You can specify more than one e-mail address by separating them with commas.
- 9 Enter the subject of the e-mail in the *Subject* field.
- 10 Enter a message for the e-mail, if any, in the *Message* field.
- 11 Click *Apply* or *OK*.

The subject and message, which you are specifying as a text strings, can contain any of the variables listed in [Table 24-3](#). These variables allow you to add details to your message about the segment or device generating the fault or event. All variables must be preceded by a percent sign (%). For example, the subject line could include the %v variable to display the severity of the alarm. You can also specify the width for the variables. %(*nnn*)X can be used to limit the length of the %X value to *nnn* characters. X represents any format specifier. For example, %(10)a is displays up to 10 characters of the Alarm ID.

Table 24-3 List of Variables

| Variable Parameter | Name | Description |
|--------------------|----------|---|
| a | Alarm ID | Identification number of the alarm as it is stored in the database. |

| Variable Parameter | Name | Description |
|--------------------|------------------------|---|
| c | Affected Class | Class of equipment that sent the alarm. This can be any portion of the network and is categorized in the database for indexing. |
| o | Affected Object Number | Identification number of the node that generated the alarm as it is stored in the database. |
| s | Alarm Summary String | Message describing the alarm. (This is the same as the status bar ticker-tape message.) |
| t | Alarm Type String | Description of the alarm. This matches the description in the Alarm Type column in the Alarm Summary window. |
| v | Severity Number | Alarm severity can be 4 = critical 3 = major 2 = minor 1 = informational All others are unknown. |
| n | Affected Object Name | Identification name of the node affected by the alarm. |
| p | Source Address | The source address of the agent that generated the alarm. |
| -h | Remove Default Header | Truncates the default header while sending an SMTP message. |

IMPORTANT: If you right-click in the body or subject area of SMTP Mail Notification, you get a context menu that contains all the variable parameters. You can select any desired variable to insert it in the message text. This variable is added at the current cursor position.

Launching an External Program

As part of editing the disposition of an alarm, you can set options to automatically launch any program on the Novell ZENworks Server Management server when an alarm is received. For example, you might want an alarm to launch a program that sends a message to the system administrator's pager.

Although ZENworks Server Management provides the capability to launch applications, the product does not provide any predefined programs. However, you can launch an NLM and run scripting routines or use third-party programs.

You can specify any necessary arguments or script variables in the Argument field. Arguments are passed directly to the program; text is not parsed, but is read as literal text strings. Variables must be preceded with a percent sign (%). The percent sign can be followed by an optional length field that limits the length to which the parameter can expand. You can also specify the width for the variables. `%(nnn)X` can be used to limit the length of the `%X` value to *nnn* characters. X represents any format specifier. For example, `%(10)a` displays up to 10 characters of the Alarm ID.

To set up automatic application launching:

- 1 Right-click the ZENworks Server Management site object in the left frame of ConsoleOne, then click *Properties*.

- 2 By default, Alarm Disposition page is selected. Click *Rule*, and then click *New* or *Edit*.
- 3 Click *Actions*, then select *Launching Application* page.
- 4 Select *Enable*.
- 5 Enter the complete path and name of the application in the *Application Name* field.
- 6 Enter any necessary execution arguments or script variables in the *Arguments* field, then click *Apply* or *OK*.

Arguments are passed directly to the program; text is not parsed, but is read as literal text strings. Variables must be preceded with a percent sign (%). The percent sign can be followed by an optional length field that limits the length to which the parameter can expand. You can also specify the width for the variables. *%(nnn)X* can be used to limit the length of the *%X* value to *nnn* characters. X represents any format specifier. For example, *%(10)a* will display the Alarm ID up to 10 characters.

Table 24-4 lists the variables you can use when launching a program:

Table 24-4 List of Variables you can use when launching a program

| Variable | Name | Description |
|----------|------------------------|---|
| a | Alarm ID | Identification number of the alarm as it is stored in the database. |
| c | Affected class | Class of equipment that sent the alarm. This can be any portion of the network and is categorized in the database for indexing. |
| o | Affected object number | Identification number of the node that generated the alarm as it is stored in the database. |
| s | Alarm summary string | Message describing the alarm. (This is the same as the status bar ticker-tape message.) |
| t | Alarm type string | Description of the alarm. This matches the description in the Alarm Type column in the Alarm Summary window. |
| n | Affected object name | Identification name of the node affected by the alarm. |
| p | Source Address | The source address of the agent that generated the alarm. |
| v | Severity number | Alarm severity can be 1 = severe 2 = major 3 = minor 4 = informational All others are unknown. |

IMPORTANT: If you right-click in the body or subject area of SMTP Mail Notification, you get a context menu that contains all the variable parameters. You can select any desired variable to insert it in the message text. This variable is added at the current cursor position.

SNMP Traps Forwarding

The Rule-Based Alarm Management System can be configured to forward an unmodified SNMP trap. The trap is automatically forwarded to the IP address of the target management station or server.

IMPORTANT: If you are specifying the port number, specify the port number after the IP address with a colon (:) preceding it. For example, if the IP address of the SNMP host server is 207.68.173.245, and the port number is 12345, specify it as 207.68.173.245:12345. By default, the port number is 162.

You can also delete one or more targets from the list.

To forward SNMP traps:

- 1 Right-click the ZENworks Server Management site object in the left frame of ConsoleOne, then click *Properties*.
- 2 By default, Alarm Disposition page is selected. Click *Rule*, and then click *New* or *Edit*.
- 3 Click *Actions*, then select *SNMP Trap Forwarding* page.
- 4 Select *Enable*.
- 5 Enter the IP address or the port number (optional) of the server to which you want to forward traps in the *SNMP Target Address* field, then click *Add*.

If you are specifying the port number, specify the port number after the IP address with a colon (:) preceding it. For example, if the IP address of the SNMP host server is 207.68.173.245, and the port number is 12345, specify as 207.68.173.245:12345.

The server is added to the List of Targets. Repeat this step for all servers you want to receive the traps.

- 6 Click *Apply* or *OK*.

Alarms Forwarding

The Rule-Based Alarm Management System can be configured to forward a processed alarm to other ZENworks Server Management servers. Specify the IP address or server name of the target management server in the Add Target dialog box and the alarm is automatically forwarded.

To add a server to the target list, select the ZENworks Server Management site and ZENworks Server Management host to which you want to forward alarms.

To remove a server from the list, select a server you want to remove and click Remove.

To forward alarms:

- 1 Right-click the ZENworks Server Management site object in the left frame of ConsoleOne, then click *Properties*.
- 2 By default, Alarm Disposition page is selected. Click *Rule*, and then click *New* or *Edit*.
- 3 Click *Actions*, then select *Alarm Forwarding* page.
- 4 Select *Enable*.
- 5 To add a target server to receive the alarms:

5a Select the Novell ZENworks Server Management site to which you want to forward alarms in the *Site Name* field.

5b Select the Novell ZENworks Server Management host to which you want to forward alarms in the *Site Host* field.

5c Click *Add*.

The server is added to the List of Targets. Repeat this step for all servers to which you want to forward alarms.

6 Click *Apply* or *OK*.

Miscellaneous

The *Miscellaneous* option on the Rule Configuration page consists of items like archiving alarms and showing received alarms on the ticker bar. In addition, it allows you to configure other options such as auto handling of an alarm and audible beeps.

The options are as follows:

- ◆ “Archiving Alarm Statistics” on page 945
- ◆ “Displaying a Ticker-Tape Message” on page 945
- ◆ “Beep On Console” on page 946
- ◆ “Auto Handling of an Alarm” on page 946
- ◆ “Displaying Alarms with Specific Severity and State” on page 947
- ◆ “Assigning Alarms to the User” on page 947

Archiving Alarm Statistics

The Rule-Based Alarm Management System system provides data to the reporting tools to generate detailed reports on alarms and network events. Enabling the *Archive* option stores the alarm in the alarm manager database on the management server.

This option is enabled by default. You should disable this option only on the types of alarms that you do not want to track and analyze. You disable this option by deselecting the *Archive* option.

To enable or disable alarm archiving:

- 1** Right-click the ZENworks Server Management site object in the left frame of ConsoleOne, then click *Properties*.
- 2** By default, Alarm Disposition page is selected. Click *Rule*, and then click *New* or *Edit*.
- 3** Click *Actions*, then select *Miscellaneous* page.
- 4** To disable alarm archiving, click the *Archive* check box to remove the check mark.
or
To enable alarm archiving, click the *Archive* check box to add the check mark.
- 5** Click *Apply* or *OK*.

Displaying a Ticker-Tape Message

The ticker-tape message is displayed in the ConsoleOne status bar and provides a summary of the most recent alarm or network event.

This option is enabled by default. You might want to edit your alarm dispositions so that only important alarms that you want to monitor display a ticker-tape message. You can disable this option by deselecting the *Show on Ticker Bar* option.

To disable or enable a ticker-tape message:

- 1 Right-click the ZENworks Server Management site object in the left frame of ConsoleOne, then click *Properties*.
- 2 By default, Alarm Disposition page is selected. Click *Rule*, and then click *New* or *Edit*.
- 3 Click *Actions*, then select *Miscellaneous* page.
- 4 To disable the ticker-tape message, click the *Show on Ticker Bar* check box to remove the check mark.
or
To enable the ticker-tape message, click the *Show on Ticker Bar* check box to add the check mark.
- 5 Click *Apply* or *OK*.

Beep On Console

The Miscellaneous page includes configuration settings such as making an audible beep in ConsoleOne. The sound alerts the user of an occurrence of an alarm. Useful applications of this function include when a server abends, when a server is downed by user, or when the file system is full.

This option is disabled by default. You should enable this option for important alarms that you want to monitor. You can enable this option by selecting *Beep on Console* option.

To enable or disable an audible beep:

- 1 Right-click the ZENworks Server Management site object in the left frame of ConsoleOne, then click *Properties*.
- 2 By default, Alarm Disposition page is selected. Click *Rule*, and then click *New* or *Edit*.
- 3 Click *Actions*, then select *Miscellaneous* page.
- 4 To enable the audible beep function, select the *Beep on Console* check box to add the check mark.
or
To disable the audible beep function, select the *Beep on Console* check box to remove the check mark.
- 5 Click *Apply* or *OK*.

Auto Handling of an Alarm

Auto handling store the alarms directly to History instead of storing them into active alarms.

This option is disabled by default. You should enable this option for important alarm that you want to auto handle. You enable this option by selecting the *Auto Handle* option.

To auto-handle an alarm:

- 1 Right-click the ZENworks Server Management site object in the left frame of ConsoleOne, then click *Properties*.

- 2 By default, Alarm Disposition page is selected. Click *Rule*, and then click *New* or *Edit*.
- 3 Click *Actions*, then select *Miscellaneous* page.
- 4 To enable the auto handling, click the *Auto Handle* check box to add the check mark.
or
To disable the auto handling, click the *Auto Handle* check box to remove the check mark.
- 5 Click *Apply* or *OK*.

Displaying Alarms with Specific Severity and State

Alarms with a specified severity and state, other than the Severity and State specified in the Alarm Template, are displayed in the Alarm History and Active Alarms View.

The Severity options are Critical, Major, Minor, Informational, and Unknown, and the State options are Operational, Degraded, Non-Operational, and Unknown.

To display the alarms with specific severity and state:

- 1 Right-click the ZENworks Server Management site object in the left frame of ConsoleOne, then click *Properties*.
- 2 By default, Alarm Disposition page is selected. Click *Rule*, and then click *New* or *Edit*.
- 3 Click *Actions*, then select *Miscellaneous* page.
- 4 Select the Severity and State from the drop-down list.
- 5 Click *Apply* or *OK*.

Assigning Alarms to the User

The users that you have added in the Manage Users list are displayed in the *Assign to User* drop-down list. You can select any user from the drop-down list to assign an alarm to that user.

To assign an alarm to a specific user:

- 1 Right-click the ZENworks Server Management site object in the left frame of ConsoleOne, then click *Properties*.
- 2 By default, Alarm Disposition page is selected. Click *Rule*, and then click *New* or *Edit*.
- 3 Click *Actions*, then select *Miscellaneous* page.
- 4 Select a user from the drop-down list.
- 5 Click *Apply* or *OK*.

For more details on managing users, refer [“Managing the Users” on page 950](#)

24.3.4 Performing Actions on Rules

You can perform the following activities on a rule associated with an alarm template:

- ♦ [“Creating and Configuring a Rule” on page 948](#)
- ♦ [“Editing a Rule” on page 949](#)
- ♦ [“Copying a Rule” on page 949](#)
- ♦ [“Deleting a Rule” on page 949](#)

- ◆ “Printing a Rule” on page 949
- ◆ “Exporting a Rule to a File” on page 950
- ◆ “Managing the Users” on page 950
- ◆ “Changing the Precedence of Rule” on page 951
- ◆ “Enabling and Disabling a Rule” on page 951

Creating and Configuring a Rule

You can create a rule to automatically perform an action when an alarm occurs.

- 1 In ConsoleOne, right-click the Novell ZENworks Server Management site object, then click *Properties*.
- 2 By default, Alarm Disposition page is selected. Click *Rule*, and then click *New*.
- 3 In Rule Configuration dialog box, specify the name of the rule. The name you specify is added to the title bar of the Rule Configuration page.
- 4 (Optional) Specify a description for the rule you are creating.
- 5 Click *Conditions*. By default, the *Source Address* is selected.
 - 5a Specify the source addresses from where you want users to get an alarms. For more details on Source Addresses, refer to “[Source Addresses](#)” on page 936
- 6 Click *Alarms*, then select *Severity*, *State* and *Alarms* as per your requirements. For more details on Alarms, refer to “[Severity, State, and Specific Alarms](#)” on page 937
- 7 Click *Time Intervals*, then select specific time durations. For more details on Time Intervals, refer to “[Time Intervals](#)” on page 939

IMPORTANT: To configure a rule, you must define at least one of the Condition.

- 8 Click *Actions*.
- 9 Click the *Enable* option. Specify *SMTP Server*, *From*, *To*, *Subject*, and *Message*. For more details on SMTP Mail Notification, refer to “[SMTP Mail Notification](#)” on page 941
- 10 Click *Launching Application* and select *Enable*. For more details on Launching Application, refer to “[Launching an External Program](#)” on page 942
- 11 Click *SNMP Trap Forwarding* and select *Enable*. For more details on SNMP Trap Forwarding, refer to “[SNMP Traps Forwarding](#)” on page 944
- 12 Click *Alarm Forwarding* and select *Enable*. For more details on Alarm Forwarding, refer to “[Alarms Forwarding](#)” on page 944
- 13 Click *Miscellaneous*, then select the options as per your requirements. For more details on the Miscellaneous options, refer to “[Miscellaneous](#)” on page 945

IMPORTANT: To configure a rule, you must define at least one of the Action.

- 14 Click *Apply*, then click *OK*.

The newly created rule is appended to the *Available Rules* list and the description that you specified in Properties page is displayed under Rule Description on the Available Rule page.

Editing a Rule

- 1 In ConsoleOne, right-click the Novell ZENworks Server Management site object, then click *Properties*.
- 2 Click the *Alarm Disposition* page.
- 3 On the Rules page, select the rule you want to edit from the *Available Rules* list, then click *Edit*.
- 4 In the Rule Configuration dialog box for that rule, modify the details as per your requirements.
- 5 Click *OK*.

Copying a Rule

You can create a copy of the any existing rule listed in the *Available Rule* list. While creating a copy of a rule, you can also modify the details of a rule.

- 1 In ConsoleOne, right-click the Novell ZENworks Server Management site object, then click *Properties*.
- 2 Click the *Alarm Disposition* page.
- 3 On the Rules page, select the rule you want to copy from the *Available Rules* list, then click *Copy*.
- 4 (Optional) In the Rule Configuration dialog box for that rule, modify the details of a rule.
- 5 Click *OK*.

Deleting a Rule

Rules displayed in the *Available Rules* list can be deleted from the list. You can delete one or more rule entries to remove it from the list.

- 1 In ConsoleOne, right-click the ZENworks Server Management site object, then click *Properties*.
- 2 Click the Alarm Disposition page.
- 3 On the Rules page, select the rule you want to delete from the *Available Rules* list, then click *Delete*.
You can select more than one rule by pressing Ctrl + clicking the rules you want to select.
- 4 Click *Yes* in Confirm Deletion dialog box.

Printing a Rule

You can also print any of the rule listed in Available Rules list.

- 1 In ConsoleOne, right-click the Novell ZENworks Server Management site object, then click *Properties*.
- 2 Click the Alarm Disposition page.
- 3 On the Rules page, select the rule you want to print from the *Available Rules* list, then click *Print*.
You can print more than one rule by pressing Ctrl + clicking the rules you want to select.
- 4 In Print Rule Summaries, click *Print*.

Exporting a Rule to a File

You can export the rule information into the following file formats:

- ◆ Comma-delimited (.csv)
- ◆ HTML (.html)
- ◆ Tab-delimited (*.txt)
- ◆ Blank-space-delimited (*.txt)

- 1 In ConsoleOne, right-click the ZENworks Server Management site object then, click *Properties*.
- 2 Click the *Alarm Disposition* page.
- 3 On the Rules page, select the rule that you want to print from the *Available Rules* list, click *Print*, then click *Export*.
- 4 In the Export dialog box, select the file type you want to use to export the rule information
- 5 Specify the filename or click *Browse* to select a filename.
- 6 Click *OK*.

Managing the Users

The Manage Users page allows you to maintain a list of users and their e-mail addresses. You can use this list to select the user to whom the alarm is to be assigned. You can perform the following activities in the Manage Users page:

- ◆ “Adding a New User” on page 950
- ◆ “Editing an Existing User” on page 950
- ◆ “Deleting an Existing User” on page 951

Adding a New User

The Manage Users page allows you create a new user to whom you can assign an alarm. Added users are listed in the *Users* list on Manage Users page.

- 1 In ConsoleOne, right-click the ZENworks Server Management site object, then click *Properties*.
- 2 Click the *Alarm Disposition* page.
- 3 On the Rules page, click *Users*.
- 4 In the Manage Users dialog box, click *New*.
- 5 In Add User dialog box, specify a username and an e-mail address.
- 6 Click *OK*.

Editing an Existing User

You can edit an existing user details, which are listed in Users list. The updated user detail is displayed in the list immediately after you update the data.

- 1 In ConsoleOne, right-click the ZENworks Server Management site object, then click *Properties*.
- 2 Click the *Alarm Disposition* page.

- 3 On the Rules page, click *Users*.
- 4 In the Manage Users dialog box, select the user you want to edit, then click *Edit*.
- 5 In the Add User dialog box, modify the username and an e-mail address.
- 6 Click *OK*.

Deleting an Existing User

Users displayed in the Users list can also be deleted from the list.

- 1 In ConsoleOne, right-click the ZENworks Server Management site object, then click *Properties*.
- 2 Click the *Alarm Disposition* page.
- 3 On the Rules page, click *Users*.
- 4 In the Manage Users dialog box, select user you want to delete.
You can select more than one user by pressing Ctrl + clicking the users you want to delete.
- 5 Click *Delete*.

Changing the Precedence of Rule

As an administrator, you can also change the precedence of the rules that are listed in the *Available Rules* list.

The rule with the highest precedence is listed at the top of the list. You can move a rule up or down in the list by clicking the up-arrow or the down-arrow.

The default rule has the lowest precedence among all the rules and you cannot modify its order in the list.

Enabling and Disabling a Rule

As an administrator, you can also enable or disable a rule that is listed in *Available Rules* list.

To enable a rule:

- 1 In ConsoleOne, right-click the Novell ZENworks Server Management site object then, click *Properties*.
- 2 Click the *Alarm Disposition* page, and then click *Rules*.
- 3 In *Enabled* column, select the option pertaining to a rule you want to enable.
- 4 Click *Apply*, and then click *OK*.

To disable a rule:

- 1 In ConsoleOne, right-click the Novell ZENworks Server Management site object then, click *Properties*.
- 2 Click the *Alarm Disposition* page, and then click *Rules*.
- 3 In *Enabled* column, deselect the option of a rule you want to disable.
- 4 Click *Apply*, then click *OK*.

IMPORTANT: Even though you cannot delete or edit the default rule, you can disable it.

24.4 Maintaining the Alarm Management System

The alarm manager database on the ZENworks Server Management server increases in size each time the Alarm Management System logs an alarm.

IMPORTANT: If you do not control the size of this database, it can increase until it fills the hard disk on the management server.

To control the size of the alarm manager database, regularly delete alarms that have been resolved or alarms that are not required for future reference or action. This deletes the instance of the alarm record from the alarm manager database and thus controls the size of the database.

You can delete alarms from the Alarm History view in Novell ConsoleOne under the View menu. For more information, see [“Deleting Alarms” on page 931](#).

24.5 Troubleshooting the Alarm Management System

When the Alarm Management System receives an unsolicited SNMP trap from an agent, it locates the appropriate alarm template for the trap-type object that is defined in the MIB of the device. If the alarm template is not available, the Alarm Management System checks the IgnoreUnknownTrap flag in the `installation_volume\installation_directory\novell\zenworks\mms\mwserver\properties\alarmmanager.properties` file. If the flag value is set to True, the alarm is ignored. If the flag value is set to False, the alarm is archived in the database as an unknown trap. If the flag value is set to Yes, the alarm is ignored. If the flag value is set to No, the alarm is archived in the database as an unknown trap. The default value of the flag is set to Yes.

For the alarm to be recognized, you need to add the MIB of the device to the MIB Pool on the management server. The MIB contains the trap definitions for traps sent from the device. If the trap-type object is undefined by the Alarm Management System, it cannot resolve the type of alarm received from the trap object identifier (OID), and the alarm is unknown. See [Chapter 26, “Using the MIB Tools,” on page 995](#) for information on compiling MIBs and adding MIBs to the MIB Pool.

If you add a new device to your network, you must add the MIB to the MIB Pool. If the SNMP agent is a proxy agent hosted on a station and the software is updated, you need to update the MIB in the MIB Pool.

Understanding Server Management

25

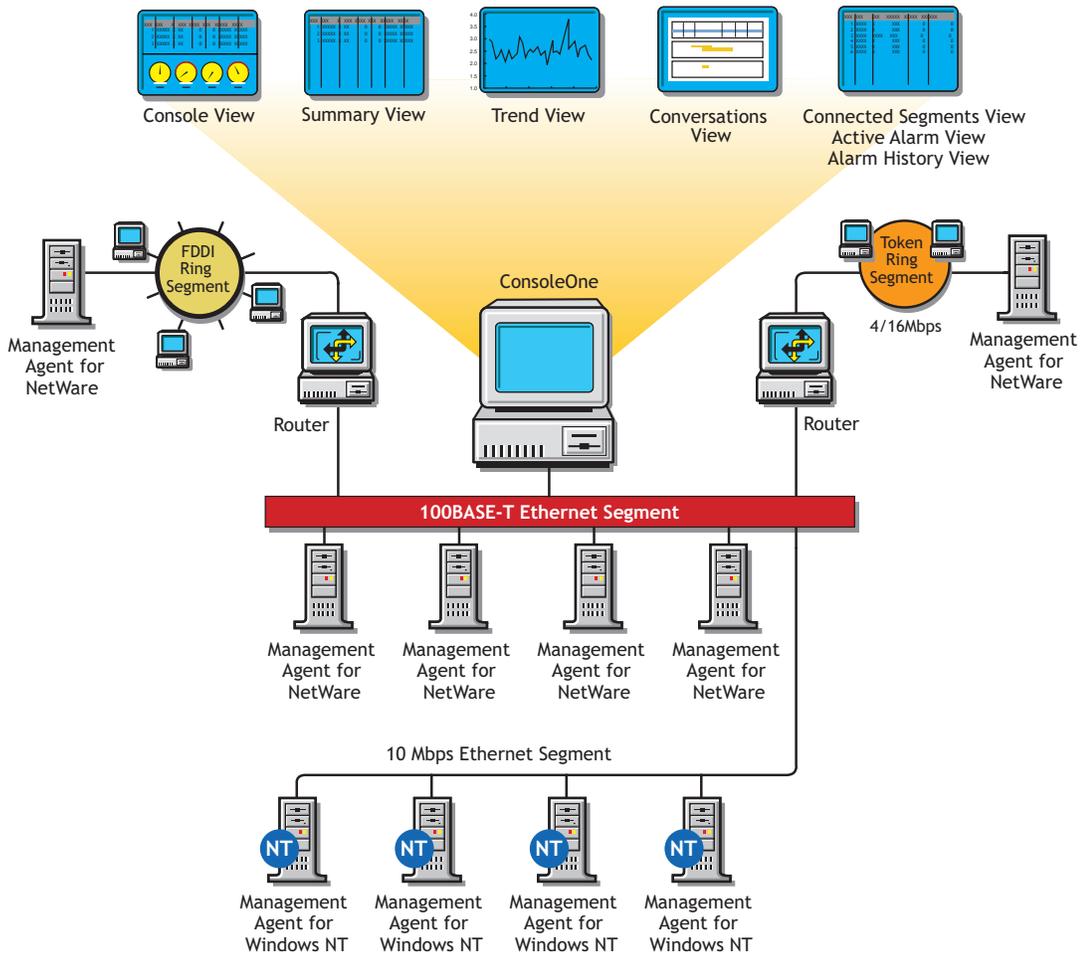
The Novell® ZENworks® Server Management components allow you to monitor, configure, and control the managed servers and nodes on your network. The SNMP-based server Management Agents for Novell NetWare® and Windows* servers provide real-time server performance data and information about server alarms and events to the network management console. By selecting a server or node from atlas page maps or hierarchical lists in the left pane of Novell ConsoleOne, you can access three main views of information:

- ◆ **Console View:** Provides details about the selected server or node. You can drill down into the server configuration to display information about the internal components of the machine, such as the devices, operating system, and services available on the machine.
- ◆ **Summary View:** Provides details about the server performance, such as alarms generated by the server, CPU utilization, and available disk space. By drilling down into the server configuration, you can also view summary information about other components, such as processors, threads, memory, and volumes.
- ◆ **Trend View:** Displays graphical representations of trend parameters, allowing you to monitor the state of a server over various periods of time. Using trend data, you can track the health status of servers, allowing you to predict potential problems and plan for future expansion of server configurations.

In addition to viewing information about the servers on your network, the server management components also enable you to configure your managed NetWare servers and execute frequently used commands from Novell ConsoleOne.

Figure 25-1 displays a functional view of the ZENworks Server Management components. It illustrates the Management Agent for NetWare and Management Agent for Windows distributed throughout a network.

Figure 25-1 Server views available from ConsoleOne



This section contains the following topics to help you understand the server management components:

- ◆ [Section 25.1, “Understanding Server Management,”](#) on page 954
- ◆ [Section 25.2, “Planning for Server Management,”](#) on page 956
- ◆ [Section 25.3, “Optimizing Server Management,”](#) on page 958
- ◆ [Section 25.4, “Managing Servers,”](#) on page 966
- ◆ [Section 25.5, “Object Hierarchy and View Details,”](#) on page 974

25.1 Understanding Server Management

The Management Agent for NetWare and the Management Agent for Windows include features that offer benefits over server management functionality included with NetWare and Windows server software.

This section includes the following topics:

- ◆ [Section 25.1.1, “SNMP-Based Server Management,”](#) on page 955
- ◆ [Section 25.1.2, “SNMP Agent Functions,”](#) on page 956

25.1.1 SNMP-Based Server Management

The main advantage of the Management Agent for NetWare and Management Agent for Windows is that they support the industry standard Simple Network Management Protocol (SNMP). SNMP is the protocol governing network management and the monitoring of network devices and their functions.

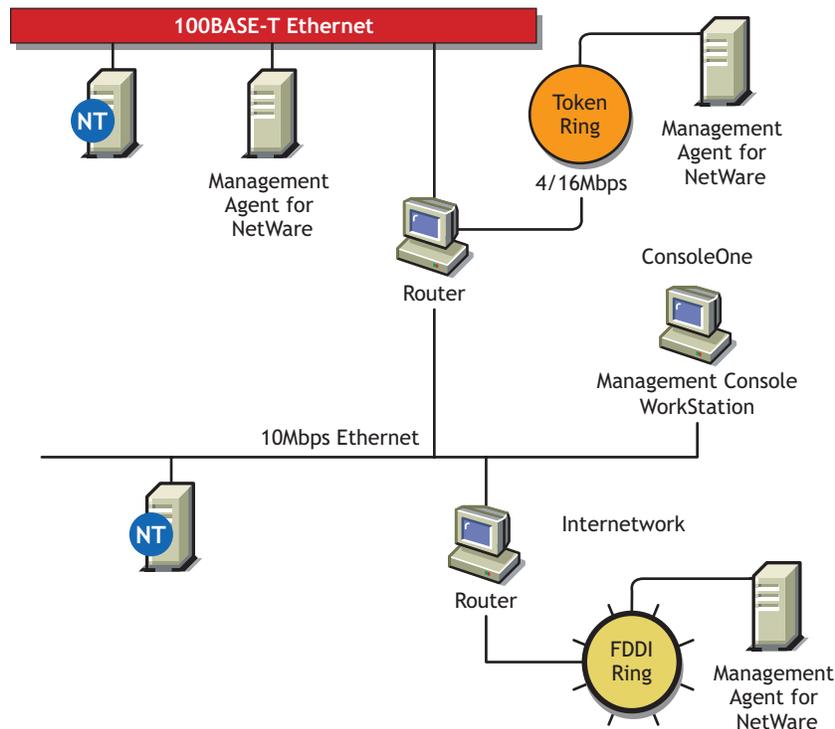
The Novell ZENworks Server Management SNMP agents support UDP/IP, IPX™, and NCP™ implementations for accepting and sending packets (datagrams). This standard mechanism allows any SNMP console or manager to request information from the Novell ZENworks Server Management SNMP agents. An SNMP console can be any console that supports SNMP; the Novell ZENworks Server Management Novell ConsoleOne fully supports SNMP v.1 communication.

SNMP Agents

The Novell ZENworks Server Management SNMP agents run on NetWare and Windows servers in your network. The agents monitor servers, collecting historical data and dynamic data in response to requests from Novell ConsoleOne. An administrator at the Novell ZENworks Server Management Novell ConsoleOne can request data simply by clicking a representative icon for any device, operating system, or service discovered on a server.

Figure 25-2 illustrates an internetwork using the Management Agent for NetWare and Management Agent for Windows and the Novell ZENworks Server Management Novell ConsoleOne.

Figure 25-2 The NetWare/Windows Agent in an Ethernet/FDDI network/token ring



25.1.2 SNMP Agent Functions

The functionality of the Management Agent for NetWare and Management Agent for Windows (the Novell SNMP-based agents for NetWare and Windows servers) can be divided into the following areas:

- ◆ Collecting Statistics
 - ◆ Monitoring: Server monitoring provides instant information about various monitored elements of the server, such as CPU utilization, memory size, cache buffers, connected users, volumes, disks, disk space usage per user, network adapters, print queues, print jobs, and loaded NetWare Loadable Module™ (NLM™) files on NetWare or Windows servers.
 - ◆ Trending: Trends provide historical data about various server objects and can be displayed in a diagram on the SNMP console. Trends are stored at the server side, which eliminates the need for continuous polling from an SNMP manager, and this data can be accessed via SNMP by any Novell ZENworks Server Management Novell ConsoleOne or other SNMP-based console.
- ◆ Alarm Notification: More than 580 different types of alarms or events (SNMP traps) can be sent from any NetWare server to the Novell ZENworks Server Management system or to any other SNMP-based console.

Any Windows system, security, or application event is converted to an SNMP trap and sent to the Novell ZENworks Server Management system or to any other SNMP-based console.

The alarms inform the administrator about events that have occurred or thresholds which have been crossed.

- ◆ Configuration Management: The Management Agent for NetWare enables network administrators to remotely configure NetWare servers. There are 187 SET parameters on the NetWare server that can be used to tune the server's performance. Administrators can view settings and change all parameters from any Novell ZENworks Server Management Novell ConsoleOne.

The SNMP agents must be installed on any server that you want to manage. For information on installing the SNMP agents, or if you have already installed the agent software to servers that you want to manage, see “[Management and Monitoring Services Installation](#)” in the *Novell ZENworks 7 Server Management Installation Guide*.

25.2 Planning for Server Management

A baseline defines the typical activity of your network servers. Keeping a baseline document of activity on a server lets you determine when the activity is atypical. To create a baseline activity, you should gather statistical information when the server is functioning typically.

This section contains the following information to help you plan your server management strategy:

- ◆ [Section 25.2.1, “Creating a Baseline of Typical Server Activity,” on page 957](#)
- ◆ [Section 25.2.2, “Using the Baseline Document,” on page 957](#)
- ◆ [Section 25.2.3, “Server Baseline Document Tips,” on page 957](#)

25.2.1 Creating a Baseline of Typical Server Activity

For server statistics such as CPU utilization, you should create a trend graph that plots information over a period of time. Statistics sampling that gathers data over a short period of time can be misleading. If you modify the server's configuration, it is useful to create another baseline against which you can compare future activity.

There are two ways to create baseline documents. The first is to create them manually by printing the various trend graphs for which you want to maintain baselines. The other way is to use the server management health reports as your baseline documents. For more information on creating and generating health reports, see [“Managing the Server Management Health Reports” on page 1151](#). In either case, the data gathered can be exported into programs, such as spreadsheets, for further analysis and to maintain records over time.

25.2.2 Using the Baseline Document

The following sections will help you plan and use the baseline document:

- ◆ [“Using Baseline Documents to Set Alarm Thresholds Appropriately” on page 957](#)
- ◆ [“Using Baseline Documents to Track Server Utilization” on page 957](#)
- ◆ [“Use Baseline Documents in Troubleshooting” on page 957](#)

Using Baseline Documents to Set Alarm Thresholds Appropriately

You should set alarm thresholds for statistics on servers monitored by the SNMP agent software, so that if the threshold is exceeded, you are notified at Novell ConsoleOne. Setting alarm threshold values for statistics on a server eliminates the need for you to constantly monitor polled server statistics for problems.

Server Management components provide default values for thresholds set on server statistics; rising and falling statistics generate an alarm when a threshold is surpassed.

Using Baseline Documents to Track Server Utilization

By comparing current server performance statistics against the performance recorded in your baseline document, you can determine how performance is affected by server configuration changes. This comparison also helps you plan for growth and justify upgrades and expansion. You can view graphs of real-time trends and historical trends over hourly, daily, weekly, monthly, and yearly periods.

Use Baseline Documents in Troubleshooting

By knowing what the typical server activity is, you can recognize atypical activity, which might help you isolate the cause of a problem.

25.2.3 Server Baseline Document Tips

You should include the following key characteristics in each server baseline document:

- ◆ [“CPU Utilization” on page 958](#)
- ◆ [“Cache Buffers” on page 958](#)

- ♦ [“File Reads and Writes” on page 958](#)
- ♦ [“Volume Utilization” on page 958](#)
- ♦ [“Running Software” on page 958](#)

CPU Utilization

The CPU Utilization statistic indicates how busy the microprocessor is. High CPU utilization can cause slow network response time. Utilization is likely to be higher at some times during the day (for example, when users log in to the network in the morning, or access e-mail), week, or month. Tracking CPU utilization helps you track the load on the server processor at peak and low times. This information helps you determine the effect of current system and application processor demands and analyze the impact on performance.

Cache Buffers

Virtually all processes are handled through server cache, a block of server memory (RAM) in which files are temporarily stored. Cache buffers greatly increase server performance and enable workstations to access data quicker because reading from and writing to memory is much faster than reading from or writing to disk. The optimum cache buffer is 65% to 75% of total server memory (more does not hinder performance). Low cache buffers can cause slow server performance and abend. Service degrades noticeably at 45% of total server memory.

File Reads and Writes

By tracking data about file reads and writes in your baseline, you might be able to determine whether a bottleneck is caused by the disk I/O channel. For example, if an increasing number of “server busy” packets are sent to users and there is also an increase in the file read and write number, the cause of the bottleneck might be a slow disk I/O channel or bad disk adapter driver.

Volume Utilization

Tracking volume utilization is primarily for capacity planning. By tracking the volume space used over time, you can accurately predict when you must purchase additional storage. Tracking volume utilization can also help you prevent the server from running out of disk space.

Running Software

By including information about running software in your baseline, it is easier to spot a problem application when comparing software on different servers. It is useful to also include the memory each application uses. Then, if the server is running short of memory, you can quickly see which applications are using the most memory.

25.3 Optimizing Server Management

Examine each of the configuration options in the sections that follow to determine whether you require any of the functionality provided:

- ♦ [Section 25.3.1, “Setting Default Trends and Thresholds,” on page 959](#)
- ♦ [Section 25.3.2, “Controlling Alarm Generation,” on page 963](#)
- ♦ [Section 25.3.3, “Defining Recipients for SNMP Alarms,” on page 965](#)

25.3.1 Setting Default Trends and Thresholds

You can modify the default trends and threshold values from within Novell ConsoleOne or manually modify files on servers that have the Management Agent for NetWare or Management Agent for Windows software installed.

When server agents are first loaded, the initial (default) values for trends and thresholds are read from the `ntrend.ini` file (NetWare) or the `n_nttren.ini` file (Windows). The initial values are also used whenever a new trend file is created. A new trend file is created when an instance of a monitored object (volume, disk, interface, and so on) is discovered on the server.

Figure 25-3 is a sample excerpt from an `ntrend.ini` file:

Figure 25-3 A sample `ntrend.ini` file

```
#-----#
# #Parameter | Sample | Trend | Threshold |
# | Interval | Buckets | Enbl | Rising | Falling | Enbl | Type |
#-----#-----#-----#-----#-----#-----#-----#
NUMBER_LOGGED_IN_USERS 5 60 1 100 90 1 rising
NUMBER_LOGGED_IN_USERS 7 8928 1 90 81 1 rising
NUMBER_CONNECTIONS 5 60 1 0 0 0 rising
NUMBER_CONNECTIONS 7 8928 1 0 0 0 rising
FILE_READS 5 60 1 0 0 0 rising
FILE_READS 7 8928 1 0 0 0 rising
FILE_WRITES 5 60 1 0 0 0 rising
FILE_WRITES 7 8928 1 0 0 0 rising
FILE_READ_KBYTES 5 60 1 0 0 0 rising
FILE_READ_KBYTES 7 8928 1 0 0 0 rising
FILE_WRITE_KBYTES 5 60 1 0 0 0 rising
FILE_WRITE_KBYTES 7 8928 1 0 0 0 rising
LSL_IN_PACKETS 5 60 1 0 0 0 rising
LSL_IN_PACKETS 7 8928 1 0 0 0 rising
LSL_OUT_PACKETS 5 60 1 0 0 0 rising
LSL_OUT_PACKETS 7 8928 1 0 0 0 rising
NCP_REQUESTS 5 60 1 0 0 0 rising
NCP_REQUESTS 7 8928 1 0 0 0 rising
CPU_UTILIZATION 5 60 1 90 81 1 rising
CPU_UTILIZATION 7 8928 1 80 72 1 rising
CACHE_BUFFERS 5 60 1 45 40 1 falling
CACHE_BUFFERS 7 8928 1 0 0 1 falling
CODE_DATA_MEMORY 5 60 1 0 0 0 rising
CODE_DATA_MEMORY 7 8928 1 0 0 0 rising
```

After the Management Agent for NetWare and Management Agent for Windows software is running, trend and threshold values can be changed (using Novell ConsoleOne) by making use of the threshold-setting features of Novell ZENworks Server Management. If the server is brought down, it retains the last trend and threshold settings that were set. Initial values are reset when any of the following situations occurs:

- ◆ Trend files have been deleted manually.
- ◆ If the server configuration is modified, for example, by adding a new volume, disk, or interface.

IMPORTANT: Trends are not maintained for CD volumes. Therefore, changing trend parameters for CD volumes has no effect.

The following sections contain information to help you modify initial trend and threshold values:

- ◆ [“Changing the Initial Trend Values” on page 960](#)
- ◆ [“Changing the Initial Threshold Values” on page 962](#)

Changing the Initial Trend Values

The trend values in the `ntrend.ini` file (NetWare) and `n_nttren.ini` file (Windows) specify the time interval (Sample Interval) at which a particular trend parameter is sampled, the duration of time for which those samples are kept (Trend Buckets), and whether this sampling parameter is enabled (Enbl). For each value specified by a line in the `ntrend.ini` file or `n_nttren.ini` file, a trend record is stored in a separate file in the `sys:\ntrend` directory on a NetWare server and the `\trenfile` directory on a Windows server.

Figure 25-4 depicts a line in the `ntrend.ini` file for the `NUMBER_LOGGED_IN_USERS` trend parameter with a Sample Interval of 5, Trend Buckets specified at 60, and the enable parameter specified at 1 (enabled).

Figure 25-4 A sample line from an `ntrend.ini` file

```
#-----#
# # Parameter | Sample | Trend | Threshold |
# # | Interval | Buckets | Enbl | Rising | Falling | Enbl | Type |
#-----#-----#-----#-----#-----#-----#
NUMBER_LOGGED_IN_USERS 5 60 1 100 90 1 rising
```

The following sections describe how to set or alter each of the parameters required for a trend file:

- ◆ [“Setting the Sample Interval” on page 960](#)
- ◆ [“Setting the Trend Buckets” on page 961](#)
- ◆ [“Enabling or Disabling a Trend File” on page 962](#)
- ◆ [“Backing Up Trend Data” on page 962](#)

You can specify more than one sampling interval or duration for any trend parameter by creating another line in the `ntrend.ini` file or `n_nttren.ini` file.

Setting the Sample Interval

The trending software enables you to collect samples of a specified parameter at any of 12 possible time intervals (Sample Interval), from 5 seconds to 1 day.

Each of these sample intervals is specified by a code number in the `ntrend.ini` file and the `n_nttren.ini` file. Table 25-1 specifies the codes used in the `ntrend.ini` and `n_nttren.ini` files for the permitted sample intervals. For example, if you want to sample a particular trend parameter once every hour, you would use the code 9.

Table 25-1 Codes to be used in the `ntrend.ini` and `n_nttren.ini` files for the permitted sample intervals

| Sample Interval | Code |
|-----------------|------|
| 5 seconds | 1 |
| 10 seconds | 2 |
| 15 seconds | 3 |
| 30 seconds | 4 |
| 1 minute | 5 |
| 5 minutes | 6 |

| Sample Interval | Code |
|-----------------|------|
| 15 minutes | 7 |
| 30 minutes | 8 |
| 1 hour | 9 |
| 4 hours | 10 |
| 8 hours | 11 |
| 1 day | 12 |

Setting the Trend Buckets

After you have determined a sample interval for collecting samples, you must set a duration of time for which you want to collect samples. For example, if you selected a sample interval of one hour for a particular parameter, you might decide that you want to be able to review the state of that parameter for every hour over the duration of a day.

You determine the duration of time for which a parameter is collected by the number of trend buckets you specify. You must specify a trend bucket for each sample that is collected over a specific period of time. For example, to review the state every hour for 1 day, 24 trend buckets (1 per hour x 24 hours in a day) are required.

The number of trend buckets required for any particular time duration and sample interval is calculated easily. However, for your convenience, [Table 25-2](#) shows the number of trend buckets required for each sample interval allowed, for each of seven possible time durations of from 1 hour to 1 year:

Table 25-2 Number of trend buckets required for each sample interval allowed

| Sample Interval | 1 Hour Duration | 1 Day Duration | 1 Week Duration | 1 Month Duration | 3 Months Duration |
|-----------------|-----------------|----------------|-----------------|------------------|-------------------|
| 5 seconds | 720 | 17280 | 120960 | 535680 | 1607040 |
| 10 seconds | 360 | 8640 | 60480 | 267840 | 803520 |
| 15 seconds | 240 | 5760 | 40320 | 178560 | 535680 |
| 30 seconds | 120 | 2880 | 20160 | 89280 | 267840 |
| 1 minute | 60 | 1440 | 10080 | 44640 | 133920 |
| 5 minutes | 12 | 288 | 2016 | 8929 | 26784 |
| 15 minutes | 4 | 96 | 672 | 2975 | 8928 |
| 30 minutes | 2 | 48 | 336 | 1488 | 4464 |
| 1 hour | 1 | 24 | 168 | 744 | 2232 |
| 4 hours | | 6 | 42 | 186 | 558 |
| 8 hours | | 3 | 21 | 93 | 279 |
| 1 Day | | 1 | 7 | 31 | 93 |

After you set the sample interval and the time duration for trend collection, you can compute the size of trend files. The number of trend buckets possible, and the approximate size in kilobytes (in parentheses), for a given sample interval and time duration are also given in [Table 25-2](#). The size of each trend bucket is 4 bytes plus 512 bytes for the header file. For example, if the sampling interval is 5 seconds for a period of 1 hour, the file size would be 720 trend buckets x 4 bytes long (rounded to the closest 4 KB boundary) plus 512 bytes for a total of 4.5 KB. There are always as many trend files as there are enabled trends.

After a particular time duration is exceeded for a file (all the trend buckets have been filled), the oldest samples are overwritten by the most recent samples. This means that the file contains the most recent duration recorded. For example, if you select a sample interval of 1 hour for a duration of 24 hours (using 24 trend buckets), the associated file contains the trend data for the last 24 hours.

Enabling or Disabling a Trend File

Each line in the `ntrend.ini` file and the `n_nttren.ini` file contains a parameter that either enables or disables the trending value to begin creating a trend file at startup. To enable the collection of data for a trend file, set this parameter to 1. To disable the collection of data for a trend file at startup, set this parameter to 0.

Backing Up Trend Data

Trend data is not automatically backed up. If you want to back up this data, you must do so manually.

Changing the Initial Threshold Values

The default threshold values in the `ntrend.ini` file and the `n_nttren.ini` file specify when a trap is generated. User-defined values are stored in the trend file header. If the parameter rises above or falls below the set threshold value, a rising or falling trap type is sent.

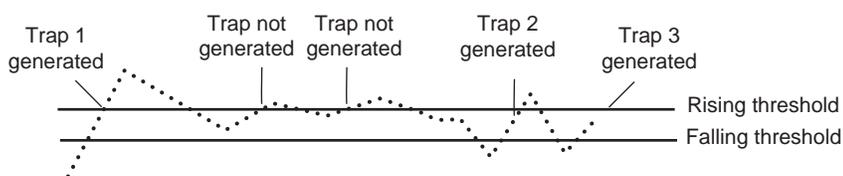
The following sections describe how to set or alter each of the parameters required for a threshold value:

- ◆ [“Setting Rising and Falling Thresholds” on page 962](#)
- ◆ [“Enabling or Disabling a Threshold Trap” on page 963](#)

Setting Rising and Falling Thresholds

Each line in the `ntrend.ini` file and the `n_nttren.ini` file contains a parameter for the rising threshold and the falling threshold. For each sample interval, a rising or falling trap can be generated as specified. After a trap is generated, another such trap is not generated until the sampled value falls below this threshold and reaches the falling threshold.

Figure 25-5 A rising threshold and falling threshold for a trap



In [Figure 25-5](#), Trap 1 is generated because it is the first time that the parameter value rises above the Rising Threshold. The next two times the parameter value rises above the Rising Threshold, a

trap is not generated because the parameter did not fall below the Falling Threshold. Trap 2 and Trap 3 are generated because the parameter value dropped below the Falling Threshold before exceeding the Rising Threshold.

Enabling or Disabling a Threshold Trap

Each line in the `ntrend.ini` file and `n_nttren.ini` file contains a parameter that enables or disables the `ntrend.nlm` software to send traps as determined by the rising and falling thresholds. This parameter is set to 1 to enable the software to send a trap for the values given, or to 0 to disable the software from sending a trap for this parameter.

25.3.2 Controlling Alarm Generation

Each managed server has files that specify which system events result in a trap. On NetWare, the `nwtrap.cfg` and `ndstrap.cfg` files are stored in the `sys:\etc` directory. On Windows, this file is `nttrap.ini`, which is stored in the `mw\ini` directory.

On NetWare, the trap configuration file is read only when `nwtrap.nlm` is loaded; therefore, any changes made to the file do not take effect until the next time you load `nwtrap.nlm` or `ndstrap.nlm`.

The `.cfg` files on NetWare contain the list of supported traps. You can modify the `.cfg` files or `nttrap.ini` file with the following:

- ◆ Types of alarms forwarded to Novell ConsoleOne
- ◆ Community strings used for sending SNMP traps
- ◆ List of traps to be disabled, using the mask keyword
- ◆ Specific alarms that you want to prevent from forwarding

The configuration file consists of keywords and their associated data (case is ignored). Each keyword must be on a line by itself (except for mask values, where they might span several lines), and must be followed by one or more lines of associated data.

You can place comments anywhere in the file, even between a keyword and its associated information. A comment starts with a number sign (`#`), and continues to the end of the line.

Figure 25-6 is an example of an `nwtrap.cfg` file:

Figure 25-6 Sample contents of the nwtrap.cfg file

```
#
#####
#NWTRAP.CFG
#
#NWTRAP Configuration File
#
#This file specifies information to be used by NWTRAP.NLM
#The file is read and the parameters set when NWTRAP is loaded. It must
#reside on volume SYS: in the directory SYS:\ETC and must be named
#NWTRAP.CFG to be found by NWTRAP. To change the parameters, first
#edit this file, then unload NWTRAP and load it again. Any changes to this
#file will not take effect until NWTRAP is next loaded. The parameters
#are specified by using a parameter keyword followed by the desired
#parameter value.
#
#####

Community
  Public
Time Interval
  10
Severity
  Warning

mask
#   "Memory: Short term alloc failed"
#   1

#   "FileSys: Directory write error (no vol)"
#   2

#   "FileSys: File write err, by server (no path)"
#   3

#   "FileSys: File write err, by user (no path)"
#   4
```

The following sections contain information to help you control alarm generation:

- ◆ [“Setting the Time Interval \(Management Agent for NetWare Only\)” on page 964](#)
- ◆ [“Configuring Alarm Severity Levels” on page 964](#)

Setting the Time Interval (Management Agent for NetWare Only)

Sometimes an alarm repeats rapidly (several times per second or per minute) with identical or nearly identical parameters. When this occurs, the second and later alarms within a time interval are usually not as interesting as the first alarm.

To prevent the network and Novell ConsoleOne from being inundated with identical alarms, you can specify a time interval to be applied to every alarm generated. During this interval, alarms that are identical to an initial alarm are discarded.

You can define the time interval in the configuration file as follows:

```
Time Interval
n
```

where n can take any value from 0 to 232 to indicate the number of seconds that must elapse before a later alarm is not discarded.

The default time interval is 10 seconds.

Configuring Alarm Severity Levels

Use the severity keyword to set a minimum alarm severity level so that traps for lesser severity alarms are not sent.

The severity levels you can set in the `nwtrap.cfg` and `nttrap.ini` files are informational, warning, recoverable, critical, and fatal. [Table 25-3](#) lists the NetWare severity level and corresponding SNMP and Novell ZENworks Server Management severity levels:

Table 25-3 NetWare severity level and corresponding SNMP and Novell ZENworks Server Management severity levels

| NetWare Severity Level | SNMP Severity Level | ZENworks Server Management Severity Level |
|------------------------|---------------------|---|
| 0 - Informational | Informational | Informational |
| 1 - Warning | Minor | Minor |
| 2 - Recoverable | Major | Major |
| 3 - Critical | Critical | Severe |
| 4 - Fatal | Fatal | Severe |
| 5 - Operation Aborted | Fatal | Severe |
| Unrecoverable | Fatal | Severe |

The default keyword is warning. Under the default, all alarms with a severity level of warning or greater are forwarded.

25.3.3 Defining Recipients for SNMP Alarms

You can configure the Management Agent for NetWare to send SNMP traps (alarms) to the Novell ZENworks Server Management server or to other management nodes.

NOTE: For setting trap destinations on Windows servers, see the documentation on the SNMP Service provided with the Microsoft Windows operating system software.

Steps for designating trap target destinations are described in the following section.

Editing the `traptarg.cfg` File Manually (Management Agent for NetWare Only)

You can configure trap recipients by manually adding them to the `traptarg.cfg` file. This is useful for sending traps to third-party management consoles other than the Novell ZENworks Server Management server.

You must add trap recipients manually by specifying their addresses in the `traptarg.cfg` file, which is located in the `sys:\etc` directory of all NetWare servers.

The `traptarg.cfg` defines the recipients of SNMP traps. You can use this file to define recipients of SNMP traps over IPX and UDP/IP. The file is fully annotated to show you how to divide the file into IPX and UDP/IP sections and how to write the IPX and IP addresses of recipients.

The `traptarg.cfg` file is read only when SNMP is loaded. In most cases, this means bringing the server down and restarting it because a variety of modules must be unloaded and reloaded as well. Thus, any changes made to the `traptarg.cfg` file do not take effect until the next time you load `nwtrap.nlm`.

IMPORTANT: The `nwalarm.mib` file imports symbols from the Host Resources MIB (RFC1514.MIB), which can also be found in `sys:novell zenworks\mms\mwserver\mibcserver\mibserverpool\mibpool.`

25.4 Managing Servers

With the Management Agent for NetWare and Management Agent for Windows software installed on your NetWare and Windows servers, respectively, you can begin collecting data, receive alarm notifications, remotely manage configuration, and generate reports for managed servers.

Server Management tasks you can perform with Novell ZENworks Server Management include:

- ◆ [Section 25.4.1, “Displaying Server Configuration Information,” on page 966](#)
- ◆ [Section 25.4.2, “Displaying Summary Data,” on page 967](#)
- ◆ [Section 25.4.3, “Viewing Trend Data,” on page 968](#)
- ◆ [Section 25.4.4, “Managing Trend Samplings,” on page 970](#)
- ◆ [Section 25.4.5, “Configuring Server Parameters,” on page 972](#)
- ◆ [Section 25.4.6, “Executing Server Commands,” on page 972](#)
- ◆ [Section 25.4.7, “Management Site Server Status,” on page 973](#)

25.4.1 Displaying Server Configuration Information

Server configuration data is organized in a hierarchical listing expanding down from the server object. You can view information about the server's configuration, memory usage, adapters, network interfaces, disks and disk controllers, volumes, queues, users, connections, open files, NLM files (NetWare), and installed software.

To display server configuration information:

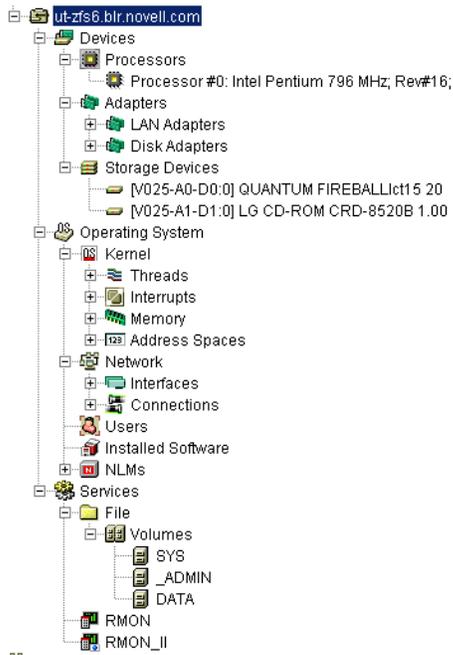
- 1 Locate the server object you want to expand.
- 2 Click the plus sign (+) next to the server object.

The server object opens in the left pane under its parent object and the server contents are displayed. Server data is grouped into the following three categories:

- ◆  Devices
- ◆  Operating System
- ◆  Services

If you are unable to view the above three categories, you must perform probe manageability on the server object. Right-click the server object then select Probe Manageability. The three categories will now be displayed.

- 3 You can drill down into the server configuration farther by clicking the plus signs next to the Devices, Operating System, and Services objects as in the following illustration:



25.4.2 Displaying Summary Data

The Summary View contains tables of statistics obtained by SNMP GET requests to the Management Agent for NetWare and Management Agent for Windows software hosted on managed servers. Statistics are updated dynamically as the server is continually polled for data. Polling utilizes SNMP GET and GET NEXT requests to update the data. You can also control the polling of a selected object by using the stop and refresh functions.

You can view summary data for server, processors, LAN adapters, disk adapters, storage devices, threads, interrupts, memory, address spaces, interfaces, connections, users, installed software, NLM files (NetWare), and volumes. For detailed information about a specific Summary view, see [Section 25.5, “Object Hierarchy and View Details,” on page 974](#).

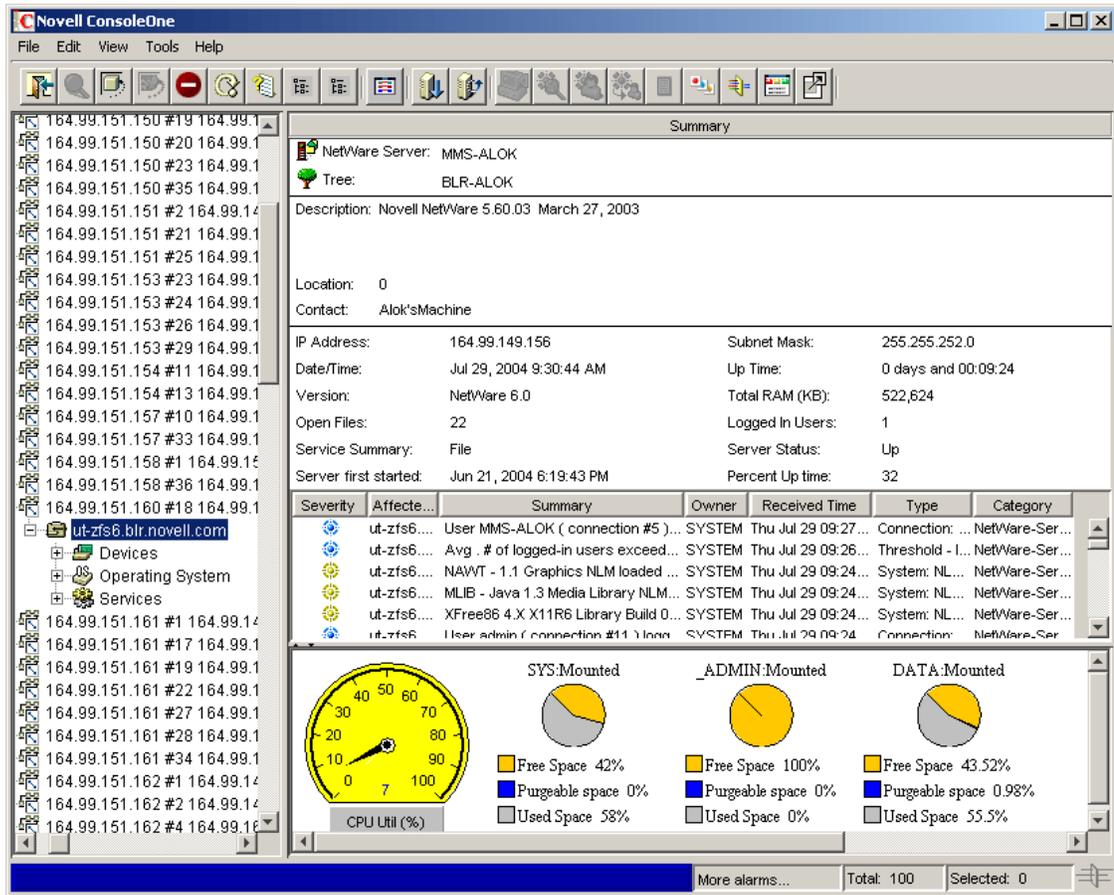
To display summary information:

- 1 Right-click the object for which you want to view summary data, then click *Views* and then click *Summary*.

The Summary View is displayed.

[Figure 25-7](#) shows the Summary View for a server object. The server summary provides descriptive information including the server's Novell eDirectory name and tree, IP address, RAM, operating system and version, IPX address, subnetmask, up time, logged-in users, open files, and status. In addition, the server summary lists all alarms, affected objects, summary, and owner and volume disk space, trend graphs for cache hits and cache buffers. The Summary view displays graphical indicators of the CPU utilization that depicts the average percentage of time that the CPU was not idle for the past minute.

Figure 25-7 Summary View for a Server Object



25.4.3 Viewing Trend Data

On a server managed by Management Agent for NetWare or Management Agent for Windows, the agents automatically gather trend data on CPU usage, memory usage, and network interface traffic. You can then view current trend data, or historical trend data by hour, day, week, month, or year from Novell ConsoleOne. In this view, the time interval that is being sampled is displayed on the x-axis. The parameter value over the sample period is plotted on the y-axis. Note that the values on the y-axis use the standard abbreviations K (for kilo), M (for mega), and G (for giga). Therefore, a value of 1K would equal 1000; similarly, a value of 1M would equal 1,000,000.

Monitoring trend data helps you with tasks such as setting trend alarm thresholds, determining who is using the server and when the server is used heavily, troubleshooting problems, balancing loads across multiple servers, and planning resources. You can also export trend view data to popular spreadsheet formats for sharing data with others.

You can view trend data for processors, LAN adapters, storage devices, memory, connections, users, and volumes. For information about a specific trend view, see [“Object View Details” on page 975](#).

To view trend statistics:

- 1 Right-click the object for which you want to view trend data then click *Views*, and then click *Trend*.

The Trend View is displayed.

The following sections describe the tasks you can perform using the Trend View:

- ♦ [“Displaying the Legend” on page 969](#)
- ♦ [“Modifying the Time Span” on page 969](#)
- ♦ [“Customizing the Trend View Display” on page 969](#)
- ♦ [“Modifying the Trend View Profile” on page 970](#)

Displaying the Legend

The Trend View legend indicates what each color in the graph represents.

To display the legend:

- 1 Click the *Legend* button  in the Trend View toolbar.

Modifying the Time Span

The Trend View time span specifies what time period the trend graphs represent. By default, a one-hour history is displayed.

To modify the time span:

- 1 Select a time span from the drop-down list in the Trend View toolbar. You can select from the following time spans:
 - ♦ 1 Hour
 - ♦ 1 Day
 - ♦ 1 Week
 - ♦ 1 Month
 - ♦ 1 Year

Customizing the Trend View Display

The Trend View provides several options for customizing the look of the screen. In customizing the view, you can choose from the following options:

- ♦ [“Displaying Grid Lines” on page 969](#)
- ♦ [“Stacking and Unstacking Graphs” on page 970](#)
- ♦ [“Scaling the Y Axis” on page 970](#)

Displaying Grid Lines

By default, the trend charts do not include grid lines.

To display horizontal and/or vertical grid lines:

- 1 To display horizontal grid lines, select the *Horizontal Grid*  button in the Trend View toolbar.
- 2 To display vertical grid lines, select the *Vertical Grid*  button in the Trend View toolbar.

Note that you remove the horizontal or vertical grid lines by clicking the same buttons.

Stacking and Unstacking Graphs

By default, all trends are displayed on a single graph with one vertical axis. However, you can customize the view so that each trend is displayed in its own separate graph.

To stack and unstack graphs:

- 1 To display the trends on separate graphs, click the *Strip Chart*  button on the Trend View toolbar.
- 2 To display trends on the same graph, click the *Stack Chart*  button on the Trend View toolbar.

Scaling the Y Axis

To display more useful information on your trend graphs, you may find that you need to modify the scale on the Y axis as follows:

- 1 To increase the scale on the Y axis, click the *Increase Y Axis*  button, which is located to the left of the graph(s).
- 2 To decrease the scale on the Y axis, click the *Decrease Y Axis*  button, which is located to the left of the graph(s).
- 3 To scale the Y axis to fit in the window, click the *Scale to Fit*  button on the Trend View toolbar.

Modifying the Trend View Profile

The Trend View profile represents the set of parameters that are displayed graphically when the Trend View is invoked. You can modify which parameters are displayed in the Trend View by editing the profile.

To edit the profile:

- 1 Click the *Profile* button  in the Trend View toolbar.

The Profile dialog box is displayed. The parameters that are currently displayed in the Trend View for the object are selected.

- 2 Edit the profile by clicking a parameter name to select or deselect it.
You can Shift+click multiple, consecutive parameters and Ctrl+click multiple, non-consecutive parameters.
- 3 Click *OK*.

25.4.4 Managing Trend Samplings

You can customize the parameters of the trend data displayed using the following options:

- ◆ [“Modifying Trend Sampling and Intervals” on page 971](#)
- ◆ [“Modifying Threshold Alarm Settings” on page 971](#)

Modifying Trend Sampling and Intervals

For each trend for which the server agents collect data, you can set sampling intervals and the number of samples stored on the server as follows:

1 Right-click the object, then click *Properties*.

2 Click the *Trend* tab.

3 Select the trend parameter you want to modify, then click *Edit*.

The Edit Trend dialog box is displayed. The trend sampling and interval settings are displayed in the Sampling Parameters section of the screen.

4 To enable or disable the sampling parameter, select the appropriate value from the *State* drop-down list.

5 To modify the time interval (Sample Interval) at which the trend parameter is sampled, select a value from the *Frequency* drop-down list.

You can select one of 12 possible time intervals from five seconds to one day.

6 Specify the duration of time for which to collect samples by entering a value in the *Number of Samples* field.

You determine the duration of time for which a parameter is collected by the number of samples (trend buckets) you specify. You must specify a trend bucket for each sample that is collected over a specific period of time. For more information on setting the number of samples required, see [“Setting the Trend Buckets” on page 961](#).

7 When you are done modifying the trend sampling and intervals, click *OK*.

Modifying Threshold Alarm Settings

You can set an alarm threshold for each trend parameter for which the Management Agent for NetWare and Windows, collects data. After you set the alarm threshold, the Management Agent for NetWare sends an alarm to Novell ConsoleOne if the trend crosses the threshold you set.

The Management Agent for NetWare tracks both rising and falling alarm thresholds. Each trend parameter has either a rising or a falling threshold associated with it; the type of threshold cannot be changed.

To change alarm thresholds through Novell ConsoleOne:

1 Right-click the object, then click *Properties*.

2 Click the *Trend* tab.

3 Select the trend parameter for which you want to modify threshold settings, then click *Edit*.

The Edit Trend dialog box is displayed. The threshold alarm settings are displayed in the Rising Alarm Parameters section of the screen.

4 To enable or disable the alarm parameter, select the appropriate value from the *State* drop-down list.

5 To set or modify the rising threshold, enter a value in the *Rising Threshold* field.

6 To set or modify the falling threshold, enter a value in the *Falling Threshold* field.

7 When you are done modifying the alarm threshold settings, click *OK*.

25.4.5 Configuring Server Parameters

In order to correct an alarm condition, fine-tune server performance, or fix other problems detected on a server, you need to modify the server configuration. Server configuration can be adjusted from Novell ConsoleOne on any NetWare server hosting the Management Agent for NetWare. SET parameters, usually set at the server console or through a remote console, can be configured from Novell ConsoleOne interface. From Novell ConsoleOne, you can see the current settings, change one or more settings, and confirm your settings before adjustments are sent to the server.

For parameter values and descriptions, see the NetWare server documentation. This information is generally found in the Utilities Reference document.

To view or modify the NetWare SET parameters from Novell ConsoleOne:

- 1 Drill down into the server you want to configure by clicking the *plus sign (+)* next to the server object.
- 2 Right-click the Operating System object, then click *Properties*.
The *Set Parameters* tab is displayed. This tab page lists the NetWare SET parameters and their current values.
- 3 Click the down-arrow icon on the *Set Parameters* tab, then click the category of SET parameters you want to display.
You can choose from the following categories: Communications, Directory Caching, Directory Services, Disk, Error Handling, File Caching, File System, Licensing Services, Locks, Memory, Miscellaneous, Multiprocessor, NCP, Service Location Protocol, Time, or Transaction Tracking.
- 4 Select the parameter you want to modify, then click *Edit*.
The Edit Parameters dialog box is displayed.
- 5 Enter the new parameter value in the appropriate field.
- 6 Indicate when you want the parameter change to take effect by selecting the appropriate radio button from the *Apply Value* box. You can choose to apply the change at the following times:
 - ◆ Now, until reboot
 - ◆ Only after reboot
 - ◆ Now, and after reboot
- 7 Click *OK*.

25.4.6 Executing Server Commands

You can execute the following frequently used NetWare server commands from Novell ConsoleOne.

- ◆ [“Loading and Unloading an NLM” on page 973](#)
- ◆ [“Mounting and Dismounting Volumes” on page 973](#)
- ◆ [“Clearing a Server Connection” on page 973](#)
- ◆ [“Restarting a Server” on page 973](#)
- ◆ [“Shutting Down a Server” on page 973](#)

Loading and Unloading an NLM

To load or unload an NLM from Novell ConsoleOne:

- 1 Right-click the NLM object, then select a command from the menu as follows:
 - ♦ To load the NLM, select *Load nlm*.
 - ♦ To unload the NLM, select *Unload nlm*.

Mounting and Dismounting Volumes

To mount or dismount a volume:

- 1 Right-click the volume object, then click *Mount Volume*.
or
Right-click the volume object then click *Dismount Volume*.

The system displays a confirmation box.

- 2 Click *OK*.

Clearing a Server Connection

You can clear a server connection when the server has crashed and left open files on the server or before bringing down the server. This is equivalent to the CLEAR STATION command that you can execute from the server console.

To clear a server connection from Novell ConsoleOne:

- 1 Locate the connection you want to close by expanding the following objects: *Server > Operating System > Network > Connections*.
- 2 Right-click the connection you want to close, then click *Clear Connection*.

Restarting a Server

To restart a server from Novell ConsoleOne:

- 1 Right-click the server object, then click *Restart Server*.

Shutting Down a Server

To shut down a server from Novell ConsoleOne:

- 1 Right-click the server object, then click *Down Server*.

25.4.7 Management Site Server Status

You can now view the status of all the Management and Monitoring Services currently running on the site server.

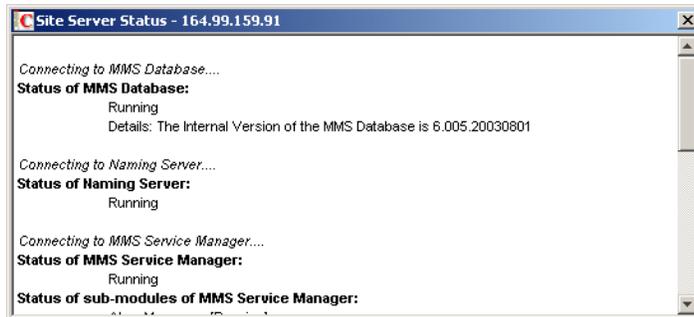
- 1 Select *Tools > Management Site Server Status*.
- 2 Specify the IP address or the DNS name of the Management Site Server.

NOTE: In case of a clustered environment, specify the IP address of a Shared volume.



3 Click *OK*.

The status of your Management Site Server is displayed in the Management Site Server window:



25.5 Object Hierarchy and View Details

When you expand a managed server object, you can view details about the contents of the server. The following sections detail the available objects on a managed server and provide information about the statistical information available in the views for each object:

- ♦ [Section 25.5.1, “Object Hierarchy,” on page 974](#)
- ♦ [Section 25.5.2, “Object View Details,” on page 975](#)

25.5.1 Object Hierarchy

[Table 25-4](#) shows the hierarchy of available objects on a managed server along with their associated icons. For more information about the available views associated with an object, follow the corresponding link.

Table 25-4 Hierarchy of available objects on a managed server along with their associated icons

| Category | Container | Sub-category | Containers | Object Containers | Objects |
|---|-----------|---|--|---|--|
|  Devices | | |  "Processors" on page 976 | |  Processor |
| | | |  "Printers" on page 993 | |  Printers |
| | | |  "Adapters" on page 979 |  LAN Adapters |  Adapter |
| | | | |  Disk Adapters |  Adapter |
| | | |  "Storage Devices" on page 977 | |  Storage Device |
| | | |  "Other Devices" on page 994 | |  Keyboard |
| | | | | |  Mouse |
|  "Ports" on page 994 | |  Parallel Port | | | |
| | |  Serial Port | | | |
|  Operating System | |  Kernel |  "Threads" on page 980 |  Thread | |
| | | |  "Interrupts" on page 981 |  Interrupt | |
| | | |  "Memory" on page 982 |  Memory | |
| | | |  "Address Spaces" on page 984 | |  Address Space |
| | | | |  "Network" on page 985 |  "Interfaces" on page 985 |
| | | | |  "Connections" on page 987 |  Connection |
| | | |  "Users" on page 988 | |  User |
| | | |  "Installed Software" on page 989 | |  Software |
|  "NLM" on page 989 | |  NLM | | | |
|  Services | |  File |  "Volumes" on page 990 |  Volume | |
| | | |  Print |  "Queues" on page 992 |  Queue |

25.5.2 Object View Details

The following sections provide details about the statistical information available in each object view:

- ♦ "Processors" on page 976
- ♦ "Storage Devices" on page 977
- ♦ "Adapters" on page 979
- ♦ "Threads" on page 980
- ♦ "Interrupts" on page 981
- ♦ "Memory" on page 982
- ♦ "Address Spaces" on page 984

- ◆ “Network” on page 985
- ◆ “Interfaces” on page 985
- ◆ “Connections” on page 987
- ◆ “Users” on page 988
- ◆ “Installed Software” on page 989
- ◆ “NLM” on page 989
- ◆ “Volumes” on page 990
- ◆ “Queues” on page 992
- ◆ “Printers” on page 993
- ◆ “Other Devices” on page 994
- ◆ “Ports” on page 994

Processors

Viewing processor speed helps you analyze and balance loads across servers. Viewing processor utilization data helps you detect problems with utilization and determine when server load is light enough to schedule tasks such as server backups. The server operating system (OS) automatically determines the CPU speed and is reported based on the OS data.

Processor speed is a major determinant of server performance. Therefore, it is important to know the processor speed of your servers when analyzing server load and balancing load across multiple servers. For example, one server might be handling twice as many users as another, but if the processor is twice as fast, the load might still be distributed correctly.

You should maintain a baseline of processor utilization for a server so that you can recognize when a server's processor utilization is higher than normal.

You can display the following views of information about the processors on your managed servers:

- ◆ “Processors Summary View” on page 976
- ◆ “Processor Summary View” on page 977
- ◆ “Processors Trend View” on page 977

Processors Summary View

You can access the **Summary View** for the Processors object container after expanding the following server objects: *Devices > Processors*. This view displays the following information for each processor object in the container:

- ◆ **Processor Number**: A unique number assigned to the processor.
- ◆ **Status**: The status of the processor is either online or offline.

The following statistics are displayed only if the processor is online:

- ◆ **Utilization %**: Processing load on this processor for the last second, expressed as a percentage.
- ◆ **Interrupts Processed**: Number of interrupts fired on this processor in the last second.
- ◆ **Time Spent in Interrupts Last Second, in Microseconds**: The amount of time in microseconds that the processor spent processing interrupts in the last second.

- ◆ **Number of Bound Threads:** The number of threads that have been bound to this processor. Threads that are bound to a processor run only on that processor. Unbound threads can be migrated from one processor to another when required.

Processor Summary View

You can select the **Summary View** for an individual processor after expanding the following server objects: *Devices > Processors > processor #x*. This view displays the following information:

- ◆ **Processor Number and Status:** A unique number assigned to the processor along with its current status. The status can be online or offline.

The following statistics are displayed only if the processor is online.

- ◆ **Utilization %:** The processing load on this processor for the last second, expressed as a percentage.
- ◆ **Interrupts Processed:** The number of interrupts fired on this processor in the last second.
- ◆ **Time Spent in Interrupts Last Second, in Microseconds:** The amount of time in microseconds that the processor spent processing interrupts in the last second.
- ◆ **Number of Bound Threads:** The number of threads that have been bound to this processor. Threads that are bound to a processor run only on that processor. Unbound threads can be migrated from one processor to another when required.

Processors Trend View

You can access the **Trend View** for the Processors object container after expanding the following server objects: *Devices > Processors*. This view displays the following graph for each processor:

- ◆ **CPU Utilization (avg. %):** The processing load on the processor for the last second, expressed as a percentage. This information is displayed only if the processor is online.

Storage Devices

You can get detailed information about the disk drives in a managed server, including disk size in megabytes, disk types, block size, and so on.

You can also view partition information for each disk drive. Partition information is especially informative because you can determine whether a partition is fault tolerant and whether the hard disk is losing data integrity.

Fault tolerance of a NetWare partition is part of the detailed information provided by Novell ZENworks Server Management. To determine whether a hard disk is losing data integrity, examine the redirected area. A number in the redirected area indicates the number of data blocks that have been redirected to the Hot Fix Redirection Area to maintain data integrity. The higher the redirected area number, the more faulty blocks there are on the hard disk. A redirected area growing over a period of time indicates a hard disk going bad.

On a NetWare server managed by the Management Agent for NetWare or a Windows server managed by the Novell ZENworks Server Management agent, the Agent automatically gathers trend data on CPU usage, memory usage, and network interface traffic. In Novell ZENworks Server Management, you can view current trend data, or historical trend data by hour, day, month, or year. Monitoring trend data helps you with tasks such as setting alarm thresholds, determining who is

using the server and when the server is used heavily, troubleshooting problems, balancing loads across multiple servers, and planning resources.

You can display the following views of information about the storage devices on your managed servers:

- ◆ “Storage Devices Summary View” on page 978
- ◆ “Storage Device Summary View” on page 978
- ◆ “Storage Devices Trend View” on page 979

Storage Devices Summary View

You can select the **Summary View** for the Storage Devices container object after expanding the following server objects: *Devices* > *Storage Devices*. This view provides the following information for each storage device on the server:

- ◆ **Disk Name**: The name of the disk drive.
- ◆ **Size (KB)**: The total size of the disk drive in kilobytes.
- ◆ **Access**: Whether the disk drive is readable and writable or just readable.
- ◆ **Status**: Whether the disk drive is operational.
- ◆ **Type**: The type of media. Media types can include hard disk, floppy disk, tape, optical disk (read-only, write once read many, and read/write), or RAM disk. If unidentifiable, other or unknown is listed in this field.
- ◆ **Driver Description**: The name of the driver used by the disk drive.
- ◆ **Block Size**: The amount of blocks used on the disk in kilobytes.
- ◆ **Heads**: the number of read/write heads on the disk drive.
- ◆ **Cylinders**: The number of cylinders on the disk drive.
- ◆ **Sectors/Track**: The number of sectors per track on the disk drive.
- ◆ **SCSI Target ID**: The target address for SCSI controllers or the unit number for other devices and the logical unit number for SCSI devices or the number zero for other devices.

Storage Device Summary View

You can display the **Summary View** for an individual storage device by expanding the following server objects: *Devices* > *Storage Devices* > *storage_device_x*. This view displays the following information:

- ◆ **Disk Name**: Name of the disk drive.
- ◆ **Logical ID**: The number assigned to a logical partition for identification.
- ◆ **Physical ID**: The number assigned to a physical partition for identification.
- ◆ **Type Partition**: The type of partition, including DOS, NetWare, and UNIX* partitions.
- ◆ **Size (KB)**: The size of the partition, in kilobytes.
- ◆ **Redirection Area**: The size of the entire Hot Fix Redirection Area.
- ◆ **Redirected Area**: The number of bad blocks Hot Fix found.
- ◆ **Reserved Area**: The number of Hot Fix redirection blocks reserved for system use.

- ◆ **Fault Tolerance**: The type of fault tolerance used. The possible fault tolerance types are duplex and mirrored. If there is no fault tolerance, this field contains the value None.

Storage Devices Trend View

You can select the Storage Devices **Trend View** after expanding the following server objects: *Devices* > *Storage Devices*. This view provides the following information:

- ◆ **File System Reads (#/min)**: Depicts the number of file system reads made per minute on multiple or single storage devices.
- ◆ **File System Writes (#/min)**: Depicts the number of file system writes made per minute on multiple or single storage devices.
- ◆ **File System Reads (KB/min)**: Depicts the number of file system reads per kilobyte volume made on multiple or single storage devices.
- ◆ **File System Writes (KB/min)**: Depicts the number of file system writes per kilobyte volume made on multiple or single storage devices.
- ◆ **Free Redirection Area (%)**: Depicts the percentage of total volume allocated to the disk redirection area.

Adapters

You can get detailed information about the network and disk adapters in a managed server, including I/O port, memory address, and interrupt configuration.

You can use this data to detect configuration problems such as the same address or interrupt is configured for two boards inside the server, or for a board and a component of the server's hardware. No two boards can use the same I/O port, memory address, and interrupt.

Problems with LAN adapters cause network problems, such as servers and workstations not being able to communicate. You can use the data collected on the LAN adapter to determine whether the frame type used by a network board is bound to a supported protocol. (A single network board might be bound to several protocols.)

You can immediately tell whether a problem is due to something as simple as using the wrong frame type on the workstation (for example, an Ethernet_II frame type on the server and the Ethernet_802.2 frame type on the workstation).

You can display the following views of information about the adapters on your managed servers:

- ◆ **“Adapters Summary View” on page 979**
- ◆ **“Adapters Trend View” on page 980**

Adapters Summary View

You can select the Adapters **Summary View** after expanding the following server objects: *Devices* > *Adapters* > *adapter_x*. This view provides the following information:

- ◆ **Description**: The type of adapter hardware. This field can include the following types of information: manufacturer, model, and version. Or, for network boards, this field may contain a short board name and the board's burned-in MAC address.
- ◆ **Type**: The type of adapter (for example, network card or disk storage).

- ◆ **Devices Attached:** The number of devices associated with an adapter (for example, the number of drives attached to the disk controller).
- ◆ **Driver Description:** Description of the driver for this adapter.
- ◆ **Version:** The version number of the driver software.
- ◆ **Interrupt Number:** The unique interrupt number used by the adapter.
- ◆ **I/O Port:** The unique I/O port block used by the adapter.
- ◆ **Memory:** The unique memory address space used by the adapter.
- ◆ **DMA:** The Direct Memory Access (DMA) Channel used by the adapter.
- ◆ **Slot:** The slot in which the adapter is installed.

Adapters Trend View

You can select the Adapters **Trend View** after expanding the following server objects: *Devices > Adapters > adapter_x*. This view provides the following graphs:

- ◆ **LSL Packets Received:** Depicts the number of LSL packets received by the adapter.
- ◆ **LSL Packets Transmitted:** Depicts the number of LSL packets transmitted by the adapter.
- ◆ **Packets Received:** Depicts the total number of packets received by the adapter.
- ◆ **Packets Transmitted:** Depicts the total number of packets transmitted by the adapter.

Threads

You can display information for all threads currently running on a managed server. A thread is recognized as an independent unit of execution.

You can display the following view of information about the threads on your managed servers:

- ◆ **“Threads Summary View” on page 980**

Threads Summary View

You can select the Threads **Summary View** after expanding the following server objects: *Operating System > Kernel > Threads*. This view provides the following information:

- ◆ **Name:** The application thread name.
- ◆ **Share Group:** The Application share groups and their associated threads and shares.
- ◆ **Parent Module:** Module (NLM) associated with this thread.
- ◆ **State:** The state of the thread, which can be one of the following: initializing, invalid, ready, running, suspended, terminated, or zombie.
- ◆ **Suspended Due To:** Reason the thread is suspended. If the thread is not in a suspended state, this field is blank.
- ◆ **Execution Time, Microseconds:** Amount of time in the last second that the processor spent executing the thread's code.
- ◆ **Stack Size, Bytes:** Size of the thread's stack.
- ◆ **Soft Affinity:** Processor on which the thread preferentially executes, but from which it can migrate when necessary.

- ◆ **Hard Affinity:** Indicates whether the thread is explicitly bound to a specified processor for the thread's lifetime. If the thread runs only on a specified processor, it is able to exploit the processor's cache state. If the thread is allowed to run on any available processor, the field value is zero.

Interrupts

You can display information for the registered interrupts on a managed server. On a multiprocessing system, interrupt information is displayed for all processors combined and individually for each online processor.

You can display the following views of information about the interrupts on your managed servers:

- ◆ [“Interrupts Summary View” on page 981](#)
- ◆ [“Interrupts Service Routines View” on page 981](#)

Interrupts Summary View

You can select the Interrupts **Summary View** after expanding the following server objects: *Operating System > Kernel > Interrupts*. This view provides the following information:

- ◆ **Name:** The name of the interrupt routine.
- ◆ **Interrupt Number:** Number for this service routine.
- ◆ **Processor:** Number of the processor.
- ◆ **Type:** The type of interrupt service routine. It can be one of the following:
 - ◆ **Bus:** A device I/O interrupt that is used (for example, by disk or LAN drivers).
 - ◆ **Local:** A hardware platform-specific interrupt local to an individual processor.
 - ◆ **System:** An interrupt category that is reserved for systems with unique interrupt requirements.
 - ◆ **Interprocessor:** An interrupt that is generated by one processor to affect another processor.
 - ◆ **Timer:** An interrupt that provides timer services for the OS as well as preemption support. (In multiprocessing systems, timer interrupts are local to a processor.)
- ◆ **Service Routines:** Number of service routines that are launched when this interrupt occurs.
- ◆ **Interrupt Occurrences:** Number of times in the last second that the interrupt occurred and was processed.
- ◆ **Execution Time:** Amount of time in the last second that the processor spent processing this interrupt.
- ◆ **Spurious Interrupts:** Number of times since the server started that an interrupt fired that should not have occurred.

Interrupts Service Routines View

The Interrupts Service Routines View provides information about the memory address spaces defined on the server.

NetWare runs in the OS address space (kernel), along with LAN drivers, storage device drivers, MONITOR, and Server Management Agents. OS address space is backed by physical memory.

All other address spaces are user space (ring 3) and are backed by virtual memory. Applications running in user space cannot cause the server to abend if the address space faults.

You can select the Service Routines View after expanding the following server objects: *Operating System > Kernel > Interrupts*. This view provides the following information:

- ◆ **Name**: The name of the interrupt service routine.
- ◆ **Service Routine Number**: Service Routine Number associated with this service routine.
- ◆ **Processor Number**: Processor number this routine is running on.
- ◆ **Interrupt Number**: Interrupt number associated with this service routine.
- ◆ **Interrupts Processed Last Second**: Number of interrupts that were processed by the ISR during the last second.

Memory

You can display the following views of information about the memory on your managed servers:

- ◆ [“Memory Summary View” on page 982](#)
- ◆ [“Memory Trend View” on page 982](#)
- ◆ [“Disk Cache View” on page 983](#)
- ◆ [“Virtual Memory View” on page 984](#)

Memory Summary View

You can select the Memory **Summary View** after expanding the following server objects: *Operating System > Kernel > Memory*. This view provides the following information:

- ◆ **Type**: The type of memory (for example, DOS, allocated memory, cache buffers, or code and data memory).
- ◆ **Unit Size (bytes)**: The size of the memory allocation.
- ◆ **Total (KB)**: The number of memory units × the unit size.
- ◆ **Units Used**: The number of memory units that have been allocated.
- ◆ **Used (KB)**: The number of KB of memory that has been allocated.

The Memory Summary View also provides a pie chart depicting memory usage on the system.

Memory Trend View

You can select the Memory **Trend View** after expanding the following server objects: *Operating System > Kernel > Memory*. This view provides the following graphs:

- ◆ **Cache Buffers (%)**: The percentage of memory allocated to cache buffers.
- ◆ **Code and Data Memory (%)**: The percentage of memory allocated to code and data.
- ◆ **Allocated memory (%)**: The amount of allocated memory.
- ◆ **Dirty Cache Buffers (%)**: The amount of dirty cache buffer memory.

Disk Cache View

This view displays utilization for disk cache memory. Use cache utilization statistics to determine when you need to install more RAM for cache. You can select this view after expanding the following server objects: *Operating System > Kernel > Memory*. It provides the following information:

- ◆ **Short Term Cache Hits %:** Percentage of requests in the last second for disk blocks that were already in cache memory. When the requested data is already in memory, disk reads don't need to be made. If this value falls below 98%, consider installing more RAM for cache. Also compare with Long Term Cache Hits.
- ◆ **Short Term Cache Dirty Hits %:** Percentage of requests in the last second for disk blocks that were already in cache memory but were dirty. Dirty cache must be written to disk before being used. Also check Long Term Dirty Cache Hits and LRU Sitting Time.
- ◆ **Long Term Cache Hits %:** Cumulative percentage of requests for disk blocks that were already in cache. When the requested data is already in memory, disk reads don't need to be made. Use this cumulative percentage to assess overall disk cache utilization. If this value falls below 90%, install more RAM for cache.
- ◆ **Long Term Cache Dirty Hits %:** Cumulative percentage of requests for disk blocks that were already in cache memory but were dirty. (Before dirty cache can be used, it must be written to disk.) Use this cumulative percentage to assess overall disk cache utilization. If this value is high or steadily incrementing, add more RAM for cache. Also check LRU Sitting Time.
- ◆ **Total Cache Blocks Allocated:** Cumulative number of requests for disk cache blocks that have been made since the server was started or rebooted. This value is the sum of the values of Allocated from Available List and Allocated from Least Recently Used (LRU). If the value of Allocated from Available is much higher, the server has sufficient RAM for cache. If the value of Allocated from LRU is high, install more RAM for cache.
- ◆ **Cache Blocks Allocated from Available List:** Number of requests for disk cache blocks that were filled by blocks in the available list (blocks that were not being used). When there are no free blocks available, requests are filled from the LRU list of cache blocks. If this value is much higher than the Allocated from LRU value, the server has sufficient RAM for cache.
- ◆ **Cache Blocks Allocated from LRU:** Number of requests for disk cache blocks that were filled by blocks from the Least Recently Used cache blocks. The system writes pending requests from the LRU cache block to disk then frees the block for the current request. Because LRU caches used only when no other cache is available, a steadily incrementing count indicates more RAM is needed.
- ◆ **Number of Times in Last 10 Minutes that the OS Had to Wait:** Number of times in the last 10 minutes that the OS waited for an LRU block in order to fulfill a request. If this value is greater than 7, install more RAM for cache.
- ◆ **Number of Times OS Had to Wait:** Number of times that the OS waited for an LRU block in order to fulfill a request.
- ◆ **Total Number of Times the Write Request Was Delayed:** Number of times a write request was delayed because there were too many writes to perform or because the disk channel was busy. A high value indicates either that the disk channel has too much I/O traffic or that you need to install more RAM for cache.

- ◆ **Number of Times the Request Was Re-tried:** Number of times a disk cache request had to be retried because the target block was being used. If this value is high or steadily incrementing, install more RAM for cache.

Virtual Memory View

This view displays information about the virtual memory system. Use these statistics to monitor the efficiency of server memory usage. If these values are fairly stable over time and if server performance is satisfactory, the server has adequate memory for its load. For example, if the value of Page faults increases, this indicates that the server performance is degrading. Conversely, if the Free swap pages value increases, it is an indication of better server performance.

You can select this view after expanding the following server objects: *Operating System > Kernel > Memory*. It provides the following information:

- ◆ **Total Page-In Requests:** Number of requests that were made to move virtual memory from swap files since the server was started (server up time).
- ◆ **Page-In Requests in Last 5 Seconds:** Number of requests to move 4 KB virtual memory pages from swap files.
- ◆ **Total Page-Out Requests:** Number of requests that were made to move virtual memory to swap files since the server was started (server up time).
- ◆ **Page-Out Requests in Last 5 Seconds:** Number of requests to move 4 KB virtual memory pages to swap files.
- ◆ **Total Swap Pages:** Number of 4 KB pages in this server's virtual memory system. (The size of the swap file in memory pages is the total number of bytes divided by 4 KB.) The size of the swap file grows or shrinks dynamically to match the memory requirements of the server's load.
- ◆ **Free Swap Pages:** Number of 4 KB pages that are available for use by the virtual memory system.
- ◆ **Reserved Swap Pages:** Number of 4 KB pages that are reserved by the virtual memory system.
- ◆ **Total Page Faults:** Number of times the virtual memory system retrieved from the swap file since the server was started (server up time).
- ◆ **Page Faults in Last 5 Seconds:** Number of times in the last five seconds that the virtual memory system retrieved from the swap file. (This means that accessed memory wasn't backed by physical memory.)

Address Spaces

NetWare runs in the OS address space (kernel) along with LAN drivers, storage device drives, MONITOR, and Server Management Agents. OS address space is backed by physical memory.

All other address spaces are user space (ring 3) and are backed by virtual memory. Applications running in user space cannot cause the server to abend if the address space faults.

You can display the following view of information about address spaces on your managed servers:

- ◆ **“Address Spaces Summary View” on page 985**

Address Spaces Summary View

You can select the Address Spaces **Summary View** after expanding the following server objects: *Operating System > Kernel > Address Spaces*. This view provides the following information:

- ◆ **Name**: Name of the virtual memory address space where this module runs.
- ◆ **Number of NLMs Loaded**: Count of NLM programs loaded in this address space. NetWare, LAN drivers, storage device drivers, MONITOR, and Server Management Agents are loaded in OS address space (kernel). A server application, such as GroupWise®, Lotus Notes*, or an Oracle* database, can be loaded in its own address space (user space or ring 3).
- ◆ **Mapped Pages**: Total number of physical memory pages backing this address space. Note that the OS address space (kernel) is the only address space backed by physical memory.
- ◆ **Restarted**: Total number of times this address space faulted and restarted automatically. A value of zero (0) indicates that no fault has occurred. A non-zero value indicates that an address space has faulted and recovered. Follow online Troubleshooting documentation for core dump instructions for address spaces.
- ◆ **Memory in Use, Bytes**: Amount of allocated memory in use.
- ◆ **Memory Not in Use, Bytes**: Amount of unused allocated memory.
- ◆ **Memory As Overhead, Bytes**: Amount of memory used for managing the allocation pool plus the amount of memory fragmentation.
- ◆ **Total Blocks**: Number of memory blocks that are in use and that are available at the request of the NLM.
- ◆ **Blocks in Use**: Number of memory blocks that were allocated and used.
- ◆ **Block Not Used**: Number of memory blocks that were allocated but not used.

Network

You can display the following view of information about the network activity on your managed server:

- ◆ **“Network Trend View” on page 985**

Network Trend View

You can access the **Trend View** for the Network object container after expanding the following server objects: *Operating System > Network*. This view displays the following graph for each network adapter:

- ◆ **Packets Received (KB/min)**: The number of kilobytes received by the adapter for the last minute.

Interfaces

You can display the following view of information about the network interfaces on your managed server:

- ◆ **“Interfaces Summary View” on page 986**
- ◆ **“Interfaces Statistics View” on page 986**

Interfaces Summary View

You can access the **Summary View** for the Network object container after expanding the following server objects: *Operating System > Network > Interfaces*.

This view displays the following information:

- ◆ **Frame Type**: The frame type that is bound to this logical board.
- ◆ **MAC Address**: The MAC address of the interface.
- ◆ **Description**: Text describing the interface board.
- ◆ **Line Speed**: The number of bits per second transmitted on this board.
- ◆ **Type**: The type of interface (for example, Ethernet CSMACD).
- ◆ **Logical Board #**: The number assigned to this logical board.
- ◆ **Logical Board Name**: The name assigned to this logical board.
- ◆ **Protocols**: The protocols to which the logical board is bound (for example, IP, ARP, or IPX).

Interfaces Statistics View

You can access the Statistics View for the Network object container after expanding the following server objects: *Operating System > Network > Interfaces*.

This view displays the following information:

- ◆ **Frame Type**: The frame type that is bound to this logical board.
- ◆ **MAC Address**: The MAC address of the interface.
- ◆ **MTU**: The size of the largest datagram which can be sent/received on the interface.
- ◆ **Admin Status**: The desired state of the interface.
- ◆ **Oper Status**: The current operational state of the interface.
- ◆ **Bytes In**: The total number of bytes received on the interface.
- ◆ **Bytes Out**: The total number of octets transmitted out of the interface.
- ◆ **Ucast Packets In**: The number of subnetwork-unicast packets delivered to a higher-layer protocol.
- ◆ **Ucast Packets Out**: The total number of packets that higher-level protocols requested be transmitted to a subnetwork-unicast address.
- ◆ **Nucast Packets In**: The number of non-unicast packets delivered to a higher-layer protocol.
- ◆ **Nucast Packets Out**: The total number of packets that higher-level protocols requested be transmitted to a non-unicast address.
- ◆ **Discards In**: The number of inbound packets which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol.
- ◆ **Discards Out**: The number of outbound packets which were chosen to be discarded even though no errors had been detected to prevent their being transmitted.
- ◆ **Errors In**: The number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.

- ◆ **Errors Out:** The number of outbound packets that could not be transmitted because of errors.
- ◆ **Unknown Protocols In:** The number of packets received via the interface which were discarded because of an unknown or unsupported protocol.

The Clear Counters button in this view resets the values only on the Management Console and not on the Server. This is done to enable the user to get the current data from the server.

Connections

You can display the following views of information about the connections on your managed server:

- ◆ [“Connections Summary View” on page 987](#)
- ◆ [“Connections Trend View” on page 988](#)
- ◆ [“Open Files View” on page 988](#)

Connections Summary View

The Connections **Summary View** displays information and statistics for the connections on the selected server. For example, this view displays the number of files currently being accessed by the server and by other clients. Certain files, such as hidden files that support Novell eDirectory, are always open. You can select this view after expanding the following server objects: *Operating System > Network > Connections > connection_x*.

This view provides the following information:

- ◆ **Connection# Login Name:** A string indicating the connection number and login name. Note that connection 0 (zero) is used by the system. The login name is the Novell eDirectory full distinguished name where applicable.
- ◆ **Client Address:**
 - IP: *xxx.xxx.xxx.xxx:port number*
 - IPX: *network:node:socket*
- ◆ **Connection Time:** The date and time the connection was established.
- ◆ **Privileges:** A connection can have one or more of the following privileges:
 - ◆ Supervisor
 - ◆ Operator
 - ◆ Auditor
 - ◆ High_Privilege
 - ◆ Second_Authentication
 - ◆ Second_High_Privilege
- ◆ **Status:** The status can be one of the following:
 - ◆ Not logged in
 - ◆ Logged in
 - ◆ Need security change
 - ◆ MacStation
 - ◆ Connection abort

- ◆ Audited
- ◆ Authenticated temporary
- ◆ Audit connection recorded
- ◆ DS audit connection recorded
- ◆ Logout in progress
- ◆ Read (bytes) : Number of bytes the connection has read since it was established.
- ◆ Written (bytes) : Number of bytes the connection has written since it was established.
- ◆ NCP Requests : Number of NCP requests the connection has made since it was established.
- ◆ Open Files : Number of files that are currently opened by the connection.
- ◆ Locked Records : Number of file records that are currently locked by the connection.

Connections Trend View

You can select the Connections **Trend View** after expanding the following server objects: *Operating System > Network > Connections > connection_x*. This view provides the following graphs:

- ◆ Connections (avg. #) : The average number of connections over the last sample interval.

Open Files View

The Connection Open Files View displays information and statistics for the connection on the server. For example, this view displays the number of files currently being accessed by the server and by other clients. Certain files, such as hidden files that support Novell eDirectory, are always open. You can select this view after expanding the following server objects: *Operating System > Network > Connections > connection_x*. This view provides the following information:

- ◆ Filename : The name of the open file, including the directory path.
- ◆ Login Name : The name of the user (if any) who opened the file. If the file was opened by the system or by an NLM, the Login Name will be a zero-length string.
- ◆ Volume Name : The physical name of the NetWare volume containing the open file.
- ◆ Directory Number : A number that uniquely identifies an open file within a NetWare volume.
- ◆ Volume ID : A number that uniquely identifies a NetWare volume. The value of this object for a particular volume has the same value as the nwVolID object for the same volume.

Users

You can display the following views of information about the users on a selected server:

- ◆ **“Users Summary View” on page 989**
- ◆ **“Users Trend View” on page 989**

Users Summary View

The Users **Summary View** provides information about the users who access the selected server. You can select this view after expanding the following server objects: *Operating System > Users*. This view provides the following information about each user:

- ◆ **Login Name**: The login name of the user.
- ◆ **Disk Usage**: The amount of disk space the user has used.
- ◆ **Last Login**: The date the user last logged in to the server.
- ◆ **Account Status**: Indicates whether the user account is valid.
- ◆ **Password**: Indicates whether the user's password is valid.
- ◆ **Real Name**: The user's Novell eDirectory real name.
- ◆ **Bad Login**: The number of failed login attempts for the user. The number 65535 displayed in this view indicates that you have exhausted the maximum number of attempts to login.
- ◆ **Bad Login Address**: The network address of the location from which the user login failed, if any.

Users Trend View

The Users **Trend View** provides information about the users who access the selected server. You can select this view after expanding the following server objects: *Operating System > Users*. This view provides the following graph:

- ◆ **Logged-In Users (avg. #)**: Depicts the average number of users logged in to the server.

Installed Software

You can display the following view of information about the software that is installed on a selected server:

- ◆ **“Installed Software Summary View” on page 989**

Installed Software Summary View

The Installed Software **Summary View** provides information about the software installed on the selected server. You can select this view after expanding the following server objects: *Operating System > Installed Software*. This view provides the following information:

- ◆ **Name**: The name of the installed software module.
- ◆ **Type**: The type of software (for example, device drivers, applications, or operating system).
- ◆ **Date Installed**: The date the software was installed.

NLM

You can display the following views of information about the NLM software on a managed server:

- ◆ **“NLM Summary View” on page 990**
- ◆ **“Resource Tag View” on page 990**

NLM Summary View

The NLM **Summary View** provides information about a selected NLM. You can select this view after expanding the following server objects: *Operating System > NLMs > nlm_x*. This view provides the following information:

- ◆ **Name**: The name of the NLM.
- ◆ **Version**: The version number of the NLM.
- ◆ **Released**: The date and time the NLM was released.
- ◆ **Memory (bytes)**: The total memory in bytes used by this NLM. This is a composite total of short term memory, semi-permanent memory, and non-movable memory, cache memory allocated by the NLM plus the sizes of the code, and data sections of this instance of an NLM.
- ◆ **Description**: A text string that describes the NLM.
- ◆ **Copyright**: The copyright string for the NLM.

Resource Tag View

You can select the NLM Resource Tag View after expanding the following server objects: *Operating System > NLMs > nlm_x*. This view provides the following information:

- ◆ **Description**: The name that the owning module assigned to this resource tag.
- ◆ **Number in Use**: The number of instances of the resource tag.
- ◆ **Resource Type**: The type of resource tag that is being tracked (for example, semaphores or processors).
- ◆ **Address Space**: Name of the address space where the module that owns the resource tag is running.

Volumes

NetWare server disk storage space is divided into volumes. You can view various data about the volumes mounted on a server, such as size, free space, how the volumes are distributed across disks, and which users are using the space. For individual volumes you can view data on configuration, open files, segments, and usage. The available views of data include:

- ◆ **“Volume Summary View”** on page 990
- ◆ **“Volume Trend View”** on page 991
- ◆ **“Open Files View”** on page 991
- ◆ **“Volume Segment View”** on page 992
- ◆ **“Volume Usage View”** on page 992

Volume Summary View

The Volume **Summary View** provides details about a single volume. You can select this view after expanding the following server objects: *Services > File > Volumes > volume_x*. This view provides the following information:

- ◆ **Size (KB)**: The size of the volume in kilobytes.

- ◆ **Free (KB)** : The amount of free space on the volume in kilobytes. As files are added or expanded, this number approaches zero. A pie chart shows you how much of the total volume size is free.
- ◆ **Used (KB)** : The amount of space, which is determined by subtracting the free disk space from the total volume size.
- ◆ **Status** : Whether the volume is mounted. If the volume is not mounted, only the volume name is listed.
- ◆ **Namespaces** : Namespaces that are supported on the volume. Namespaces supported are DOS, Macintosh*, NFS*, FTAM, and OS/2*.
- ◆ **Attributes** : Attributes of the volume. Possible attributes are block sub-allocation, file compression, data migration, auditing, and read-only. A volume can have a combination of attributes, such as read-only volume with block sub-allocation.
- ◆ **# Logical Segment** : The number of segments comprising this volume.
- ◆ **DS Name** : The volume's full Directory Services distinguished name or a zero-length string if not applicable.
- ◆ **Non-Purgable** : The amount of space (in kilobytes) taken by the deleted files whose purge dates have not yet expired. Non-purgable space can be reclaimed as free space when the deleted files become eligible to be purged.
- ◆ **Block Size** : The block size on the volume in bytes.
- ◆ **Dir Slots** : The total number of directory table entries available on the volume.
- ◆ **Used Dir Slots** : The number of directory table entries that are currently in use.
- ◆ **File System Name** : The type of file system on the volume is either remote or local. The File System Name value is listed only if the volume is remote. In this case, the file system name is the remote mount point; for example, SITE1:/usr/x.

Volume Trend View

You can select the Volume **Trend View** after expanding the following server objects: *Services > File > Volumes > volume_x*. This view provides the following graph:

- ◆ **Volume % Free Space** : The percentage of space still available on the volume.

Open Files View

The Volume Open Files View displays a table of all open files on the volume. If it is opened by more than one connection, multiple entries for the same file will appear in the table. You can select the Open Files View after expanding the following server objects: *Services > File > Volumes > volume_x*. This view provides the following information:

- ◆ **Filename** : The name of the open file, including the directory path.
- ◆ **Connection #** : The number of the connection that opened the file.
- ◆ **Login Name** : The name of the user (if any) who opened the file. If the file was opened by the system or by an NLM, the Login Name will be a zero-length string.
- ◆ **Directory Number** : A number that uniquely identifies an open file within a NetWare volume.
- ◆ **Volume ID** : A number that uniquely identifies a NetWare volume.

Volume Segment View

The Volume Segment View provides information about the segments on a volume. You can select this view after expanding the following server objects: *Services > File > Volume > volume_x*. As long as the Volume Segment View is displayed, the server is polled for data and the view is constantly updated with real-time information. This view provides the following information about each segment on the selected volume:

- ◆ **ID:** The number assigned to the volume segment for identification.
- ◆ **Logical Partition ID:** The number assigned to a logical partition for identification.
- ◆ **Physical Partition ID:** The number assigned to a physical partition for identification.
- ◆ **Size:** The size of the segment.
- ◆ **Fault Tolerance:** The type of fault tolerance used on the segment. Possible types are duplex and mirrored. If there is no fault tolerance, the value is None.
- ◆ **Disk Drive:** The name of the disk drive on which the segment resides.

Volume Usage View

The Volume Usage View provides information about the amount of volume space in use per user. As long as the Volume Usage View is displayed, the server is polled for data and the view is constantly updated with real-time information. You can select this view after expanding the following server objects: *Services > File > Volumes > volume_x*. This view provides the following information per volume user:

- ◆ **Used KB:** Number of kilobytes currently in use.
- ◆ **Limit KB:** Number of kilobytes to which a user is limited.
- ◆ **Username:** The user's login name.

Queues

You can display the following views of information about the queues on a managed server:

- ◆ [“Queues Summary View” on page 992](#)
- ◆ [“Queue Summary View” on page 993](#)
- ◆ [“Queue Trend View” on page 993](#)

Queues Summary View

The Queues **Summary View** provides the following information about the print queues on the managed server:

- ◆ **Queue Name:** The name of the queue.
- ◆ **Type:** The type of queue (for example, archive queue, job queue, or print queue).
- ◆ **# Jobs:** The number of print jobs in the queue currently.
- ◆ **# Print Servers:** The number of print servers serviced by the queue.
- ◆ **Volume:** The volume where the queue resides.
- ◆ **Add Job State:** Indicates whether or not the queue can add jobs.
- ◆ **Attach State:** Indicates whether or not the queue can attach.

Queue Summary View

The Queue **Summary View** provides the following information about the print jobs in the selected queue:

- ◆ **Job #:** A unique number assigned to the print job.
- ◆ **Position:** The print job's order in the print queue.
- ◆ **Bytes:** The number of bytes to be printed.
- ◆ **Description:** A description of the print job.
- ◆ **User:** The username of the user who submitted the job.
- ◆ **Entry Time:** The time the job was added to the queue.
- ◆ **Control Flags:** A value representing the control flags for the job. For example, some possible control flags are service auto start, execute, user hold, or operator hold.
- ◆ **Target Time:** The date and time the job is to be printed.
- ◆ **Target Server:** The target server for the job.
- ◆ **Actual Server:** The name of the server currently processing the job.

Queue Trend View

The Queues **Trend View** provides the following graph for each queue on the managed server:

- ◆ **Wait Time of Next Ready Job (sec):** The average length of time the next job waits in the queue.

Printers

You can get the detailed information about the printers installed in a managed server, including printer name, port, driver and description, status, error conditions, etc. You can display the following views of information about the printers on your managed servers:

- ◆ **“Printer Console View” on page 993**
- ◆ **“Printer Summary View” on page 993**

Printer Console View

You can access the Console View for the Printers object container after expanding the following server objects: *Devices > Printers*. This view displays the following information for each printer object in the container:

- ◆ **Printer Name:** Name of the printer

Printer Summary View

You can display the Summary View for an individual printer by expanding the following server objects: *Devices > Printer > printer_x*. This view displays the following information:

- ◆ **Printer Name:** The name of the printer
- ◆ **Printer Status:** The current status of this printer device. The status can be idle, printing, warm-up, or unknown state.

- ◆ Error Condition: The error conditions include lowPaper, noPaper, lowToner, noToner, doorOpen, jammed, offline, or serviceRequested.

Other Devices

From this view, you can get other devices like the keyboard and the mouse installed on a managed server.

- ◆ [“Other Devices on Console View” on page 994](#)

Other Devices on Console View

This displays other devices like the keyboard and the mouse.

The information about the keyboard includes:

- ◆ Keyboard Name
- ◆ Keyboard Type
- ◆ Driver Name
- ◆ Class
- ◆ Bus Type

The information about the mouse includes:

- ◆ Mouse Name
- ◆ Mouse Type
- ◆ Driver Name
- ◆ Class
- ◆ Bus type

Ports

From this view, you can install serial ports and parallel ports on a managed server.

- ◆ [“Ports Console View” on page 994](#)

Ports Console View

This displays information about the serial ports such as COM ports and parallel ports such as LPT ports. The information about the ports includes:

- ◆ Port Name
- ◆ Controller
- ◆ Bus Type

Novell ZENworks[®] Server Management provides the tools to manage Simple Network Management Protocol (SNMP)-manageable devices on your network. This section describes the Management Information Base (MIB) tools, the SNMP MIB Compiler and the SNMP MIB Browser. It also explains how to set up and use the tools. See the following sections for more information:

- ♦ [Section 26.1, “Understanding MIB Tools,” on page 995](#)
- ♦ [Section 26.2, “Configuring MIBs and Setting Up MIB Tools,” on page 1004](#)
- ♦ [Section 26.3, “Using the MIB Browser,” on page 1007](#)
- ♦ [Section 26.4, “Maintaining MIBs,” on page 1016](#)

26.1 Understanding MIB Tools

The following sections provide information about the tasks required for managing SNMP devices using the MIB Compiler and the MIB Browser.

- ♦ [Section 26.1.1, “About MIBs,” on page 995](#)
- ♦ [Section 26.1.2, “Understanding the SNMP MIB Compiler,” on page 995](#)
- ♦ [Section 26.1.3, “Understanding the SNMP MIB Browser,” on page 997](#)
- ♦ [Section 26.1.4, “Managing Devices with MIB Tools,” on page 999](#)
- ♦ [Section 26.1.5, “Trap Definitions,” on page 999](#)

26.1.1 About MIBs

To manage a device, you must obtain a copy of the MIB or MIBs that the device supports. A MIB is an ASCII text file, written in a precise format that describes the management information available on a particular class of devices. If, for example, you have an XYZ router from company X and you want to use Novell ZENworks Server Management for managing the router, company X must provide you with the XYZ router MIB. Novell ZENworks Server Management provides many standard and vendor-proprietary MIBs, which are found in the MIB Pool folder in the MIB Server Pool folder. By default, Novell ZENworks Server Management compiles the most generally applicable of these MIBs.

If you want to compile any new MIBs, you must store them in the MIB Pool folder in the MIB Server Pool folder. The console user can select or remove MIB files from the MIB Pool folder in the MIB Server Pool folder. The MIB Compiler compiles the files listed in the MIB Pool folder in the MIB Server Pool folder.

26.1.2 Understanding the SNMP MIB Compiler

The MIB Compiler does the following:

- ♦ Parses a set of predefined SNMP MIB files written in ASN.1 and SNMP V1, V2 syntax and verifies their syntax.

- ◆ Stores the compiled files in the Novell ZENworks Server Management database, which lets all users access these compiled files from a central location.

From the console, you can easily compile and maintain the MIB files located in the MIB Server Pool. You can add or remove MIB files from the MIB Pool.

- ◆ Updates trap definitions in the alarm template database.

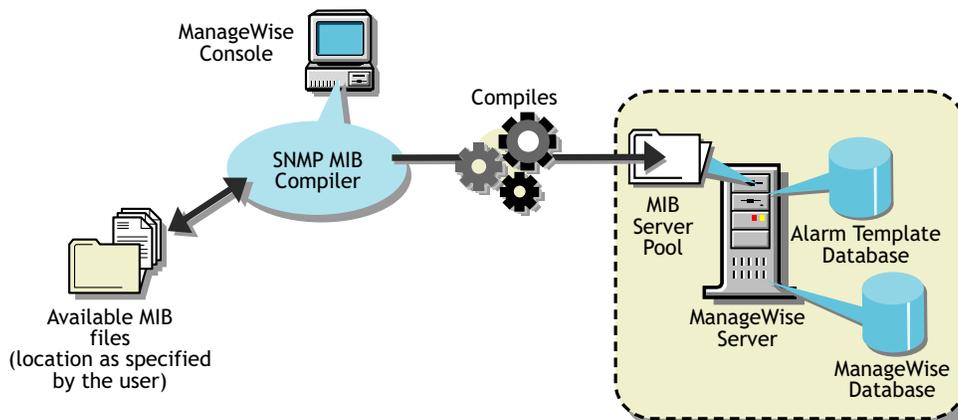
The MIB Compiler lets you introduce new SNMP alarm templates into Novell ZENworks Server Management so they can be recognized and interpreted as alarms when they arrive at the console.

The Alarm Management System interprets the annotations to trap definitions in a MIB to set the severity level and device status assigned to an alarm. The MIB files included with Novell ZENworks Server Management are already properly annotated.

- ◆ The MIB Pool on the site server now includes RFCs, which are extensively used by third-party vendors.
- ◆ The MIB Pool now includes Novell proprietary MIBs like Novell GroupWise, Novell BorderManager, Novell eDirectory, Novell Gateways, and Novell VPN.
- ◆ The SNMP Compiler Results window now includes online help on the error codes generated during compilation. The help provides you all the troubleshooting information for the error codes, to enable you to solve the problems immediately.
- ◆ The SNMP Compiler Results window now highlights the error text in Red to enable better readability.
- ◆ The MIB Compiler now ignores minor errors in the MIB definitions. This enables the MIB Compiler to compile more MIBs.
- ◆ The type, category, generator type (TCG) of any MIB need not be unique. The SNMP MIB Compiler now captures all the traps for which the TCGs are identical.
- ◆ The SNMP MIB Compiler now works with the Alarm Management System to maintain the alarm information in synchronization with the available trap information in the MIB Pool. The SNMP MIB Compiler is now faster.

Figure 26-1 demonstrates how the MIB Compiler incorporates information from the MIB files into the Novell ZENworks Server Management database:

Figure 26-1 MIB file information incorporated into the Server Management database



During installation of Novell ZENworks Server Management, the MIB files that are precompiled using the MIB Compiler are also installed. The MIB for any SNMP node you want to manage must

be compiled with Novell ZENworks Server Management. You can also integrate third-party MIBs. If you obtain a MIB file from a third-party vendor or any MIB file that was not installed with Novell ZENworks Server Management, you must compile the file using the MIB Compiler.

Using Role-Based Services with the MIB Compiler

Novell ZENworks Server Management role-based services let you assign various roles to users on your network. If your role is assigned the Enable MIB Compiler task, you can use the MIB Compiler.

See [Section 21.3, “Role-Based Administration,” on page 830](#) for more information about the role-based administration provided by Novell ZENworks Server Management.

26.1.3 Understanding the SNMP MIB Browser

The MIB Browser lets you manage SNMP-instrumented devices on the network.

To use this tool, you must have knowledge of SNMP and a good understanding of the structure of MIBs. Using the MIB Browser, you can manage nodes on the network by setting values of the MIB objects at the target nodes.

If you are familiar with the structure of an SNMP MIB, you can use the MIB Browser to retrieve data from SNMP-manageable node.

The MIB Browser lets you communicate with devices through an SNMP agent on the network over the User Datagram Protocol (UDP) or the Internet Protocol (IP). The results of SNMP commands are displayed in the MIB Browser window.

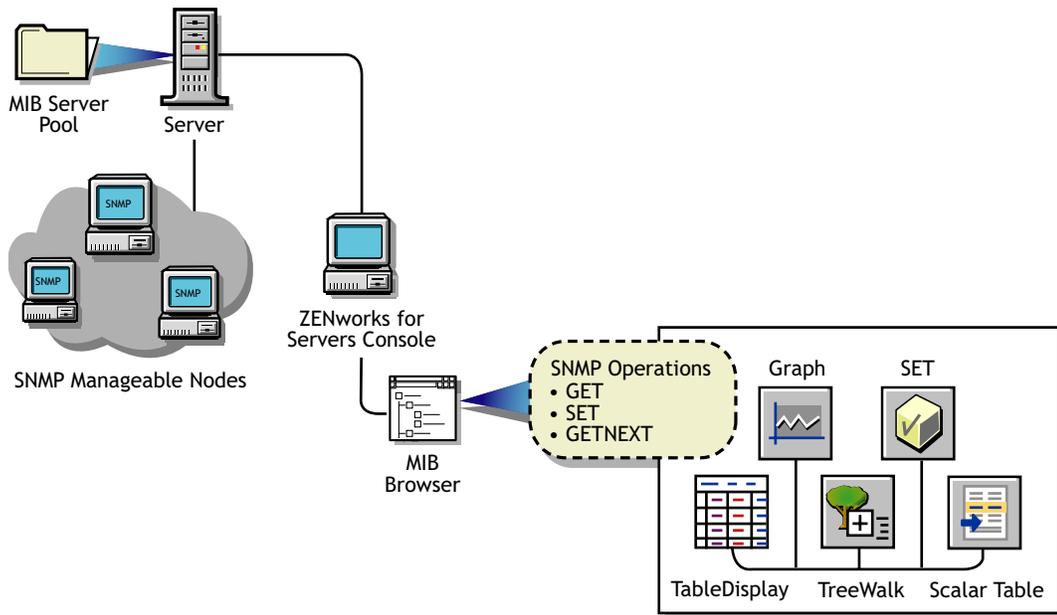
An SNMP agent is a program that provides access to management data about a particular network device and responds to SNMP Manager requests for the data. The NetWare[®] Management Agent software is an example of an SNMP agent that resides on a NetWare server. An SNMP agent resides in each manageable device on the network.

Although many Novell ZENworks Server Management windows display data retrieved from SNMP-manageable nodes, some administrators prefer the capability the MIB Browser provides for specifying the type of data they want to retrieve. Additionally, by using the MIB Browser, you can obtain some SNMP data that is not displayed in Novell ZENworks Server Management windows.

The MIB Browser takes the compiled MIB and displays the objects in a tree format. The MIB Browser also lets you walk the tree and look for the definitions of the selected MIB objects. You can set the community string to be used in the conversation between Novell ConsoleOne[®] and the SNMP-manageable node to manage the device.

[Figure 26-2](#) demonstrates the functionality of the MIB Browser:

Figure 26-2 MIB Browser functionality



The MIB Browser does the following:

- ◆ Represents the MIB information as a tree.

You can browse the objects in the MIB tree, which displays the composite OID (object identifier) for all compiled MIBs. The OID is the sequence of integers labeling each object on the path from the root of the tree to every object on the branches. The OID also describes the location of the object in the tree. For example, the novell(23) object in the tree is described as 1.3.6.1.4.1.23. For more information on the MIB tree, see [“Browsing the MIB Tree” on page 1007](#).

- ◆ Retrieves specific information about the node using the SNMP GET and GETNEXT commands.

The MIB information is displayed as:

- ◆ A table display for tabular objects

You can add new rows to the table and issue SNMP SET commands to update the columnar values of the table. For more information, see [“Modifying Instances of an SNMP Table” on page 1010](#).

- ◆ A graph display

If you choose to plot the SNMP requests, the Graph window displays the polled data of one or more MIB objects. For more information, see [“Graphing SNMP Request Results” on page 1014](#).

- ◆ A scalar table display

You can form a scalar table by combining scalar objects. You can modify the scalar entries of the table. For more information, see [“Forming Tables of Scalar Objects” on page 1013](#).

- ◆ A TreeWalk display

You can browse the OID values of scalar and tabular objects. For more information, see [“Viewing the Values of an Object and Its Child Nodes” on page 1009](#).

- ◆ Changes the information at the target node using the SNMP SET command.
You can retrieve or change the value of MIB objects if the community strings match at the target node. The node should also allow remote setting of its variables.
- ◆ Creates a profile by saving the properties of the table, scalar table, or graph.
You open the profile to view a table, scalar table, or graph of different SNMP-manageable nodes on the segment with the properties specified in the profile. For more information, see [“Using a Profile for Tables and Graphs” on page 1016](#).
- ◆ New icons for better readability including Scalar Objects, Table and the Table Entry Objects, Columnar Objects, and the MAXACCESS value.
- ◆ The Search utility now enables you to locate MIB variables in the MIB tree. You can locate them based on the MIB variable or the MIB OID.

For more information on the MIB Browser, see [Section 26.3, “Using the MIB Browser,” on page 1007](#).

Using Role-Based Services with the MIB Browser

Novell ZENworks Server Management role-based services let you assign various roles to users on your network. If your role is assigned the Enable MIB Browser task, you can use the MIB Browser.

See [Section 21.3, “Role-Based Administration,” on page 830](#) for more information about the role-based administration provided by Novell ZENworks Server Management.

26.1.4 Managing Devices with MIB Tools

Novell ZENworks Server Management lets you manage any SNMP-manageable devices on the network. In particular, you can do the following:

- ◆ Set alarm templates for receiving alarms, often referred to as SNMP traps, for these devices
- ◆ Use the MIB Browser to display and set values on these devices

Before using the MIB Browser to manage the devices, you need to perform the following tasks:

1. Acquire the necessary MIBs.
2. Add trap annotations, if required.
3. Add or remove MIBs using the [MIB Compiler](#).
4. Run the MIB Compiler to compile the MIBs in Novell ZENworks Server Management.

ASN.1 and SNMP V2 Support

The MIB Compiler supports all MIB files written in ASN.1 and SNMP V1, V2 syntax. The MIB Compiler allows relaxation of ASN.1 syntax.

26.1.5 Trap Definitions

Some SNMP MIBs define the traps that a device can send to Novell ConsoleOne when an unusual event occurs on the network. When you compile a MIB containing traps, information about those traps is added to the Novell ZENworks Server Management alarm database. When Novell ZENworks Server Management receives a trap, the information in the alarm database is retrieved

and used by Novell ZENworks Server Management to generate the alarm summary string and to determine the alarm type, alarm severity, state of the affected device, and other details.

You can improve the presentation of the alarm information in Novell ZENworks Server Management by adding annotations to the trap definitions in the MIB files. These annotations are added as comments to the trap definitions so that the MIB compiles with third-party MIB compilers.

All Novell® MIBs are annotated. If you choose not to annotate the traps in other MIBs, Novell ZENworks Server Management displays the alarms; however, they are less readable. SNMP MIBs use the TRAP-TYPE macro to define traps.

This section covers the following topics:

- ◆ “Keywords for Trap Definitions” on page 1000
- ◆ “Template Database” on page 1001
- ◆ “Keywords for Trap Annotations” on page 1001
- ◆ “Example Trap Definitions” on page 1002
- ◆ “Displaying Annotated Traps in Novell ZENworks Server Management” on page 1003
- ◆ “Formatting the SUMMARY String” on page 1003

Keywords for Trap Definitions

Table 26-1 explains a trap definition:

Table 26-1 Keywords for Trap Definitions

| Keyword | Example | Explanation |
|-------------|--|--|
| TRAP-TYPE | duplpxNetAddr | Specifies the name of the trap. For example, duplpxNetAddr represents a duplicated IPX network address. |
| ENTERPRISE | Novell NetWare-GA-alert-mib | Contains the OBJECT identifier of a node in the vendor's tree, which, together with the trap number (the 8 following the ::= in DESCRIPTION) uniquely identifies the trap. |
| VARIABLES | (osName, osLoc, tiTrapTime, tiEventValue, tiEventSeverity, tiServer) | Defines an ordered sequence of MIB objects that are passed as parameters of the trap to provide additional information about the event. For example, osName is a text string specifying the name of the server sending the trap; osLOC is a text string specifying the location of the server; tiTrapTime is an integer specifying the time the event occurred. |
| DESCRIPTION | "Two servers use the same IPX Internet address." | Provides a textual description of the semantics of the trap. |
| Trap_number | ::=8 | Defines the trap. |

Template Database

The MIB Compiler populates the alarm template database with the trap definitions in the MIB files. Any traps from the agents are stored in the database.

Keywords for Trap Annotations

Figure 26-2 lists and explains the keywords you can use to annotate traps:

Table 26-2 *Keywords for Trap Annotations*

| Keyword | Explanation |
|--------------|--|
| --#TYPE | Short name for the alarm. The name can contain a maximum of 40 characters. If this annotation is not present, the SNMP trap name is used. Every trap should have a unique type. |
| --#SUMMARY | Description of the alarm with placeholders and formatting information for the actual parameters passed with the alarm. See "Formatting the SUMMARY String" on page 1003 for more information. Without this annotation, the alarm summary string lists each SNMP parameter name followed by its value. |
| --#ARGUMENTS | List of parameters to substitute in the SUMMARY string. Parameters are substituted in the order in which they appear in the list. Each element of the list is the index (zero-based) of the parameter in the VARIABLES clause. |
| --#SEVERITY | Default severity assigned to the trap. This can be one of the following: <ul style="list-style-type: none">◆ INFORMATIONAL◆ MINOR◆ MAJOR◆ CRITICAL◆ UNKNOWN Without this annotation, the severity is displayed as UNKNOWN. |
| --#TIMEINDEX | Index of the variable in the VARIABLES clause. This index contains the time when the alarm was generated. The time is expected to be an integer representing the number of seconds since 1970 (UNIX* time). If such a variable does not exist in the VARIABLES clause, use an index greater than the total number of variables in the VARIABLE clause. |
| --#HELP | This index contains name of the help file. |
| --#HELPTAG | The index contains the reference to the Help ID of the help file that is specified in the HELP index. |

| Keyword | Explanation |
|----------|---|
| --#STATE | <p>Default state of the object when the alarm was generated. This can be one of the following:</p> <ul style="list-style-type: none"> ◆ OPERATIONAL ◆ NONOPERATIONAL ◆ DEGRADED ◆ UNKNOWN <p>Without this annotation, the state is UNKNOWN.</p> |

Note the following rules about adding trap annotations:

- ◆ Each annotation must be embedded in a comment. Everything from the double hyphen to the end of the line is treated as a comment.
- ◆ Each annotation must be on a separate line.
- ◆ Annotations must appear in the order in which they are discussed in [“Trap Definitions” on page 999](#).
- ◆ All annotations must be inserted after the DESCRIPTION clause and before the ::= clause.
- ◆ STATE and SEVERITY values are written to the alarm database the first time the MIB is compiled. If you want to modify the STATE and SEVERITY values for the alarm templates, modify these values in the corresponding MIB files and recompile using the MIB compiler.

Example Trap Definitions

The following sections explain a trap description in an SNMP trap before and after annotation:

- ◆ [“Example Trap Definition Before Annotation” on page 1002](#)
- ◆ [“Example Trap Definition After Annotation” on page 1002](#)

Example Trap Definition Before Annotation

```
dupIPXNetAddr TRAP-TYPE
ENTERPRISE Novell NetWare-GA-alert-mib
VARIABLES{osName, osLoc, tiTrapTime, tiEventValue, tiEventSeverity, tiServer}
DESCRIPTION "Two servers use the same IPX internetwork address."
::=8
```

Example Trap Definition After Annotation

```
dupIPXNetAddr TRAP-TYPE
ENTERPRISE Novell NetWare-GA-alert-mib
VARIABLES{osName, osLoc, tiTrapTime, tiEventValue, tiEventSeverity, tiServer}
DESCRIPTION "Two servers use the same IPX internetwork address."
```

-- Trap annotations are as follows:

```
--#TYPE "Duplicate IPX address"
--#SUMMARY "%s at %s and %s are using the same IPX address"
--#ARGUMENTS {0,1,5}
--#SEVERITY CRITICAL
--#TIMEINDEX 2
--#HELP "MYHELP.HLP"
--#HELPTAG 60004
--#STATE DEGRADED
::=8
```

Displaying Annotated Traps in Novell ZENworks Server Management

Assume that the dupIpxNetAddr trap shown in “[Keywords for Trap Definitions](#)” on page 1000 was received by Novell ZENworks Server Management with the following variables:

- ◆ osName = SJM-JACK
- ◆ osLoc = JACK's CORNER
- ◆ tiTrapTime = ~700000000
- ◆ tiServer = SJM-TIM

To display a trap, use the Active Alarm, Alarm History, or Alarm Detail window. The following example shows the result:

Receive Time:03/04/99 09:15:45

Alarm Type: Duplicate IPX address

Summary: SJM-JACK at JACK's Corner and SJM-TIM are using the same IPX address

Severity: Severe

State: Degraded

Formatting the SUMMARY String

The SUMMARY keyword in the trap annotation lets you provide the actual wording of the alarm summary. This wording is used by Novell ZENworks Server Management when the alarm occurs.

Placeholders within the string are replaced by actual parameters of the trap before the string is displayed by Novell ZENworks Server Management. Each placeholder format string begins with a percentage sign (%) and tells Novell ZENworks Server Management how to format the parameter that will be substituted for the placeholder in the final string.

The placeholder format strings are substituted, in order, by the parameters specified in the ARGUMENTS keyword. The ARGUMENTS keyword lists the (zero-based) index of each trap parameter as specified in the VARIABLES clause. The indexes are listed in the order in which you want them to be substituted in the SUMMARY string.

Novell ZENworks Server Management can display a maximum of 140 characters in the SUMMARY string. Use the characters to display the most relevant information about the alarm. If you have a long SUMMARY string and want to keep the line length of the MIB file reasonable, you can insert multiple, consecutive SUMMARY annotations and the strings will be concatenated. For example, the following annotations below yield the same string:

–#SUMMARY “%s at %s and %s are using the same”

–#SUMMARY “IPX address”

–#SUMMARY “%s at %s”

–#SUMMARY “and %s are”

–#SUMMARY “using the same IPX address”

Table 26-3 lists the format strings and parameter types:

Table 26-3 *Format strings and parameter types*

| Parameter Type | Format String | Printed Form |
|-------------------|---------------|---|
| BOOLEAN | %s | True or False. |
| | %d | 1 or 0. |
| INTEGER | %x | HEX. |
| | %d | DECIMAL. |
| | %t | Prints the integer or a date and time (Greenwich Mean Time). The integer represents seconds since 1970. |
| OCTET STRING | %s | Prints the text string with all control characters taken out. |
| | %m | Prints the first 6 bytes of data as a hyphen-separated MAC address. For example, 00-00-07-00-07. |
| | %x | Prints the octet string in hexadecimal. For example, 0000070007. |
| NULL | %d | Prints the number 0. |
| | %s | Prints the string NULL. |
| OBJECT IDENTIFIER | %s | Prints dot-separated decimal values. For example, 1.3.6.5.4. |
| IP Address | %s | Prints dot-separated IP address. For example, 13.56.56.56. |
| | %x | Prints a long hexadecimal value. |
| BIT STRING | %s | Prints each byte as decimal. |

26.2 Configuring MIBs and Setting Up MIB Tools

This section describes the procedural tasks for configuring MIBs and setting up the community strings for SNMP operations on an individual node. After you complete these tasks, you can perform SNMP operations using MIB Tools.

This section covers the following topics:

- ♦ [Section 26.2.1, “Annotating Third-Party MIBs for Integration with Novell ZENworks Server Management,” on page 1005](#)
- ♦ [Section 26.2.2, “Compiling MIBs for SNMP-Manageable Nodes,” on page 1006](#)

26.2.1 Annotating Third-Party MIBs for Integration with Novell ZENworks Server Management

When you compile a MIB containing SNMP traps (alarms), information about those traps is added to the Novell ZENworks Server Management alarm database. This information can then be displayed in Novell ConsoleOne.

All Novell MIBs are annotated so that the alarm information displayed in Novell ConsoleOne is easily readable. This alarm information includes a summary describing the alarm, the alarm severity, and the state of the affected node. Third-party MIB files do not necessarily contain this same information. Therefore, the information about the traps in third-party MIBs is not as meaningful when displayed in Novell ConsoleOne.

You can add annotations to third-party MIB files for the trap definitions so that the alarm information displayed in Novell ZENworks Server Management for those traps is more readable than if you compile the MIB as is. Any annotations you add to a third-party MIB are added as comments to the trap definitions. This ensures that the MIB still compiles with third-party MIB compilers.

If you do not annotate the traps in third-party MIBs, Novell ZENworks Server Management will display the alarms. The MIB Compiler displays warnings in the status display about the missing annotations.

To add annotations to a third-party MIB:

- 1 Open the MIB in a text editor.
- 2 Add any of the annotations shown in [“Keywords for Trap Annotations” on page 1001](#), by following these rules:
 - ♦ Enter annotations only between the DESCRIPTION and the “::=” clause.
 - ♦ Each annotation must be on a separate line.
 - ♦ Annotations must be in the order shown in [“Keywords for Trap Annotations” on page 1001](#).
 - ♦ Embed each annotation as a comment. Precede each annotation with two hyphens and a pound sign (#).
For example: `--#Type "type_description"`
For a full example, see [“Example Trap Definitions” on page 1002](#).
- 3 When you finish annotating trap definitions, save your changes and exit the text file.
Compile the MIB, as described in [“Compiling MIBs for SNMP-Manageable Nodes” on page 1006](#).

Use Novell ConsoleOne Alarm Disposition table to view the values for the alarm severity level and alarm state from the default values in the SNMP MIBs. If you change the value for an alarm's

severity or state after you compile the MIB, you must recompile the MIB for those changes to overwrite any changes made through the Alarm Disposition table.

26.2.2 Compiling MIBs for SNMP-Manageable Nodes

The MIB Compiler lets you manage the MIB Server Pool and also compile the .MIB files contained in the MIB Server Pool. The information in the compiled files is placed in the database on the Novell ZENworks Server Management server. The MIB Browser and the SNMP protocol decoder use this database.

The MIB Compiler also adds or updates any trap definitions to the alarm template database for use by the Novell ZENworks Server Management Alarm Management System.

The MIB Server Pool contains the list of MIB files. You can add or remove the MIB files from the MIB Server Pool.

To compile the MIBs:

1 In Novell ConsoleOne, click the Novell ZENworks Server Management server node.

2 Right-click the node then click *Properties*, and then click the *MIB Pool* tab.

The current MIB Pool lists the compiled MIB files present in the database.

3 Choose your options.

- ◆ To add MIBs, click *Add* to locate the .mib files and add them to the MIB Pool list.

The added MIBs are displayed in the adjacent list box.

When you add MIBs, you choose to integrate or exclude the trap information while compiling MIBs. If you do not integrate traps with the MIBs, only the MIB information is stored in the database on successful compilation of the MIBs. Click *Advanced*, then select the *Trap Integration* check box to integrate the trap information with the MIBs.

- ◆ To remove files from the MIB Pool list, select the MIB from the list, then click *Remove*.
- ◆ To compile the MIBs with less strict adherence to ASN.1 syntax, click *Advanced*, then select the *ASN.1 Syntax Relaxation* option.

4 Click *Compile*.

The MIB Compiler compiles all files in the MIB Pool list with the .MIB extension and updates the database. The compilation process is begun by launching a Results dialog box. This dialog box displays the status information of the MIBs including the MIBs that were successfully compiled, MIBs that were not compiled and the corresponding error message, and the status of updating the database with the MIB compile information, and the status of updating the Alarm database.

IMPORTANT: You cannot closed the Results dialog box during compilation. The *Close* button in the Results dialog box is disabled during compilation. You can close this dialog box only after the compilation is successful or failed.

5 Click *Close*.

IMPORTANT: If the SNMP MIB is not set up correctly, or an imported Request for Comments (RFC) is not available during compilation of the MIB, or any other .mib file is not available, an error message is generated in the MIB Compiler window. Add the required RFC or the dependent MIB and compile.

26.3 Using the MIB Browser

This section acquaints you with using the MIB Browser to manage SNMP-manageable nodes.

This section includes the following topics:

- ◆ [Section 26.3.1, “Browsing the MIB Tree,” on page 1007](#)
- ◆ [Section 26.3.2, “Viewing the Values of an Object and Its Child Nodes,” on page 1009](#)
- ◆ [Section 26.3.3, “Configuring a Node by Setting Object Values,” on page 1010](#)
- ◆ [Section 26.3.4, “Modifying SNMP Preferences,” on page 1010](#)
- ◆ [Section 26.3.5, “Modifying Instances of an SNMP Table,” on page 1010](#)
- ◆ [Section 26.3.6, “Forming Tables of Scalar Objects,” on page 1013](#)
- ◆ [Section 26.3.7, “Graphing SNMP Request Results,” on page 1014](#)
- ◆ [Section 26.3.8, “Using a Profile for Tables and Graphs,” on page 1016](#)

26.3.1 Browsing the MIB Tree

The MIB Browser lets you select the objects you want to display, and it sends SNMP queries to the node to obtain the data objects that you requested. It also allows SNMP operations such as GET, GETNEXT, and SET requests on a particular object in the MIB of an SNMP-managed node.

The MIB Browser periodically polls the node and continually updates the display. You can view and modify scalar and tabular data objects.

MIB Tree Browser

Within the MIB browser, the MIB Tree Browser is a graphical display of management data that consists of numerous objects.

The MIB Browser displays a composite OID tree for all compiled MIBs. Analogous to a file system, the MIB Browser shows leaf objects, which are the SNMP data objects.

The MIB Browser spans the selected node with its subtree and leaf objects and displays the name of the objects in the MIB Tree Browser. You browse from the highest level of the tree and view the leaf object values.

The top pane of the MIB Tree Browser displays the tree with the selected object. Each object is displayed as a file folder icon, followed by its SNMP name with the SubId appended in parentheses. If the object is a non-leaf node, the MIB Tree Browser also displays its children.

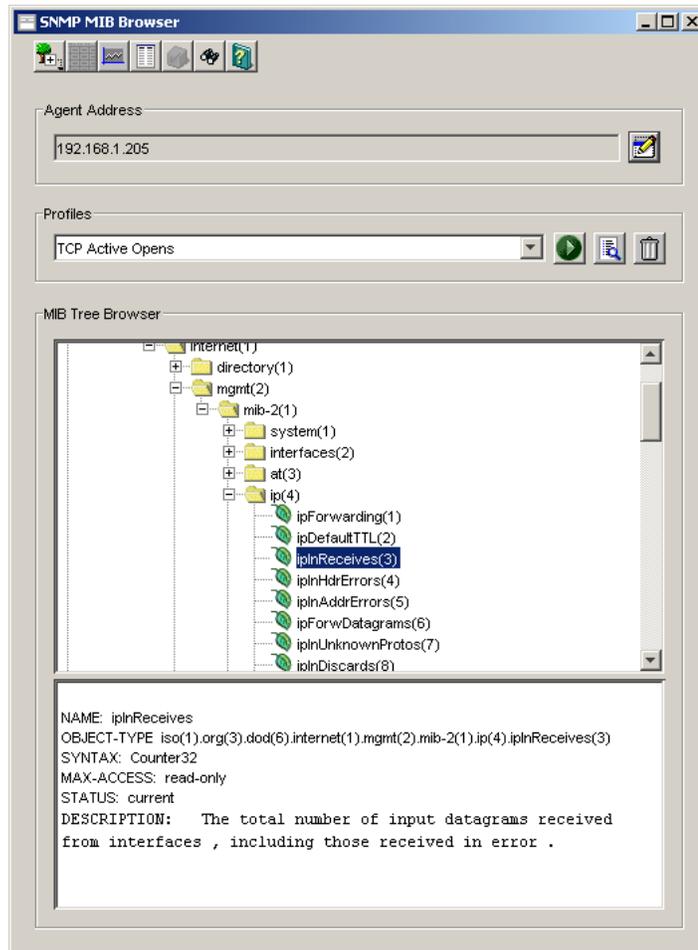
The bottom pane describes the selected object. The description is derived from the compiled MIB file. The format of the description is as follows: textual description of the object, full numeric OID and object name, ASN.1 type, size, textual convention, access, Index clause taken from the Entry object, status, and description.

For example, for an internal node SYSTEM with child nodes, the child nodes describe the properties of the SYSTEM. The OID of SYSTEM is iso(1).org(3).dod(6).internet(1).mgmt(2).mib-2(1).system(1). Another equivalent representation of this OID is 1.3.6.1.2.1.1. Note that the parent node does not have information, and the child nodes contain the properties.

The child nodes of SYSTEM are sysDescr OID(1.3.6.1.2.1.1.1), sysObjectID OID(1.3.6.1.2.1.1.2), sysUpTime OID(1.3.6.1.2.1.1.3), sysContact OID(1.3.6.1.2.1.1.4), sysName OID(1.3.6.1.2.1.1.5), sysLocation OID(1.3.6.1.2.1.1.6), and sysServices OID(1.3.6.1.2.1.1.7).

Figure 26-3 shows the MIB Browser window:

Figure 26-3 SNMP MIB Browser



To browse the MIB objects:

- 1 In Novell ConsoleOne, click the target SNMP-manageable node.
- 2 Click *File > Action > MIB Browser*.
- 3 Click the object whose values you want to view from the MIB Tree Browser.
 - ♦ To select an object, click the name text or the icon in the MIB Browser tree.
 - ♦ To expand or collapse the next level in the tree display, double-click the object.

26.3.2 Viewing the Values of an Object and Its Child Nodes

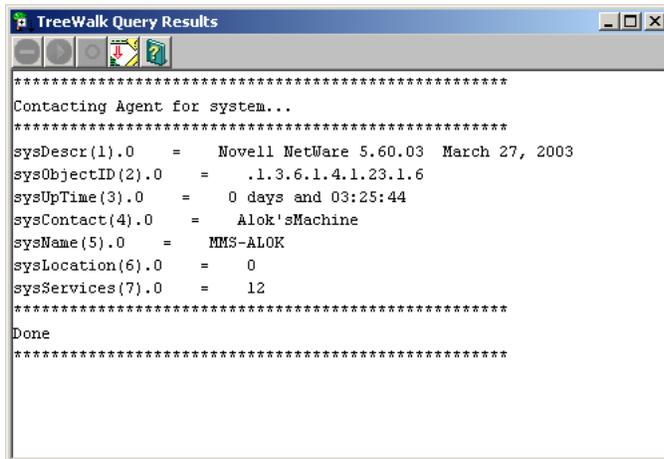
The MIB Browser spans the selected node with its subtree and leaf objects and displays its values in the TreeWalk Query Results window. You can browse the OID values of scalar and tabular objects.

To view the values of the instances of a MIB object:

- 1 In Novell ConsoleOne, click the target SNMP-manageable node.
- 2 Click *File > Action > MIB Browser*.
- 3 Click the object, then *Perform TreeWalk* for the node button.

Figure 26-4 shows the TreeWalk Query Results window:

Figure 26-4 TreeWalk Query results



If you select a leaf object, you can view the values for each instance of this object. For non-leaf objects, this window will display all the values of the child node of this object. For example, if you want to view the values of the child nodes for the object *system*, click the parent object *system*.

The display process in the Treewalk Query Results window continues recursively for all the non-leaf objects of the selected object. You can pause and resume this display in the window.

Customizing the Display of TreeWalk Query Results Window

The TreeWalk Query Results window displays the number of lines based on the settings specified in the `treewalk.properties` file.

This file, located in `\novell consoleone\version\bin\saved-views\generic` directory, contains the following setting:

```
MaximumNumberOfLine=number_of_lines_for_display
```

where *number_of_lines_for_display* is the number of lines that will be displayed at a time. The default setting is 10,000 lines. You can modify this setting. The settings will apply only if you restart Novell ConsoleOne and bring up the TreeWalk Query Results window. To clear the display in the TreeWalk Query Results window when the text buffer is full, click the *Clear* button. There may be some out of memory problems if you specify a large line setting in the `treewalk.properties` file.

26.3.3 Configuring a Node by Setting Object Values

Using the MIB Browser, you can issue an SNMP SET command to change information at an SNMP-manageable node if you have the appropriate privileges. You select a scalar object from the MIB Browser and set its value.

You can modify the values for an integer, enumerated integer, object identifier, string, and IP address object types.

To issue an SNMP SET command for a scalar object:

- 1 In Novell ConsoleOne, click the target SNMP-manageable node.
- 2 Click *File > Action > MIB Browser*.
- 3 Click a scalar object whose values you want to view, then click *Display Data As a Scalar Table*.
- 4 Specify the object value for the scalar object.
- 5 Click *OK*.

To modify columnar values of an SNMP table, see [“Modifying Instances of an SNMP Table” on page 1010](#).

26.3.4 Modifying SNMP Preferences

SNMP parameters are used to communicate with the target device. The MIB Browser lets you change the SNMP community strings or specify the transport address of a new target device.

Any SNMP operation requires these values to be set. After starting an SNMP operation, such as polling a table, changing the SNMP preferences does not affect the operation.

You can modify the following parameters:

- ♦ **Agent Address:** You can specify the IP or internal IPX address and the Domain Name System (DNS) name of the SNMP-manageable node to which you want to send an SNMP request. This node should have an SNMP agent.
- ♦ **SET and GET Community Strings:** The community string that Novell ZENworks Server Management uses must match the one expected by the SNMP agent in the managed node or the SNMP operations will fail. If the SNMP agent on the node expects a community string for SET and GET operations that is different from public (the default), you can specify the expected community string to override the default community or those community strings you set previously. You can use Unicode* or International characters for the community string.

To modify the SNMP preferences:

- 1 In Novell ConsoleOne, click the target SNMP-manageable node.
- 2 Click *File > Action > SNMP MIB Browser*.
- 3 Click *Modify SNMP Preferences*.
- 4 Specify the parameters, then click *Close*.

26.3.5 Modifying Instances of an SNMP Table

A table in an SNMP MIB is an SNMP construct derived from the structure of the MIB. Each row in the table corresponds to a row in the SNMP table.

The MIB Browser provides the Table Display window to display tabular objects you select. This window displays one or more rows from an SNMP table in a two-dimensional grid and follows the SNMP index order to display rows.

The table shows each column in the SNMP table as columns. Each column heading is derived from the SNMP table columns. The Table Display window displays the columns with their values as single or multiple rows for the MIB you selected.

SNMP allows operations on individual table entries only. The OID identifies the column and row.

From the MIB Browser, you can perform the following operations:

- ◆ Add rows to an SNMP table

For more information, see [“Adding Rows to an SNMP Table” on page 1012](#).

- ◆ Modify a row of an editable table

For more information about adding or modifying rows, see [“Adding Rows to an SNMP Table” on page 1012](#).

- ◆ Save the table as a profile

For more information about saving a table as a profile, see [“Using a Profile for Tables and Graphs” on page 1016](#).

Viewing the SNMP Table

To view the SNMP table:

- 1 In Novell ConsoleOne, click the target SNMP-manageable node.
- 2 Click *File > Action > SNMP MIB Browser*.
- 3 Click a tabular object whose values you want to view, then *Display Data As a Table*.

From the Table Display window, you can add rows or modify the rows of the SNMP table and input values for each column. For more information about adding or modifying rows of an SNMP table, see [“Adding Rows to an SNMP Table” on page 1012](#).

[Figure 26-5](#) shows the MIB Browser Table Display window.

Figure 26-5 MIB Browser table

| hostTopNReport | hostTopNIndex | hostTopNAddress | hostTopNRate |
|----------------|---------------|-------------------|--------------|
| 6,300 | 1 | FF FF FF FF FF | 497 |
| 6,300 | 2 | 00 90 21 5A BC 27 | 270 |
| 6,300 | 3 | 00 03 47 1F 3D 2F | 219 |
| 6,300 | 4 | 01 80 C2 00 00 00 | 49 |
| 6,300 | 5 | 01 00 5E 00 01 16 | 29 |
| 6,300 | 6 | 01 00 5E 7F FF FD | 26 |
| 6,300 | 7 | 01 00 81 00 01 00 | 10 |
| 6,300 | 8 | 01 00 81 00 01 01 | 10 |
| 6,300 | 9 | 01 00 1D 00 00 00 | 10 |
| 6,300 | 10 | 03 00 00 00 00 01 | 6 |
| 6,300 | 11 | 00 B0 D0 22 C3 72 | 4 |
| 6,300 | 12 | 00 00 1D 29 19 FF | 1 |
| 6,300 | 13 | 00 08 C7 C9 88 C2 | 1 |
| 6,300 | 14 | 01 00 0C CC CC CC | 1 |
| 6,300 | 15 | 00 C0 4F 37 C6 93 | 1 |
| 6,300 | 16 | 00 E0 2B 00 00 00 | 1 |
| 6,300 | 17 | 00 0F 20 33 46 59 | 1 |
| 6,300 | 18 | 00 01 03 21 A7 8D | 0 |
| 6,300 | 19 | 00 0B CD 10 C5 86 | 0 |
| 6,300 | 20 | 08 00 20 9A 0E 5E | 0 |

The MIB Browser periodically sends SNMP queries to the node to obtain the data objects you request. When you provide new values for writable objects, the MIB Browser writes these values to the node. The MIB Browser periodically polls the node and continually updates the display. You can change the polling interval by suspending the SNMP interaction or by canceling the SNMP interaction.

Adding Rows to an SNMP Table

When you add a new row to an SNMP Table, the MIB Browser generates the SNMP SET request.

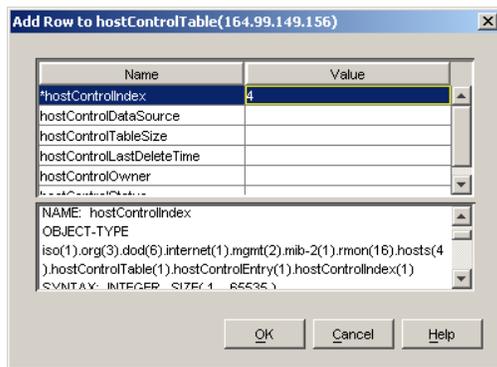
Before generating the SNMP Set request, the MIB Browser sends a GET command to the node that you selected in the MIB Browser table and retrieves the value of the object. On adding rows with the specified values for the objects, the MIB Browser issues multiple SNMP SET commands to update the SNMP table.

To add a row to an SNMP table:

- 1 Click the table object from the MIB Browser window, then *Add a New Row to the Table*.

For more information about selecting the table object, see [“Modifying Instances of an SNMP Table” on page 1010](#).

The following figure shows the Add Row to Table window:



- 2 Double-click the row.
- 3 Modify the value then click *OK*, and then click *OK*.

To add rows in an SNMP table, you must input the values for all the index rows, which are denoted by asterisks.

To modify a row of an editable table:

- 1 Open the Table window.
- 2 Click the row whose values you want to modify, then click *Issue SNMP Set request for a column* button.
- 3 Double-click the row.
- 4 Modify the value of the object then click *OK*, and then click *OK*.

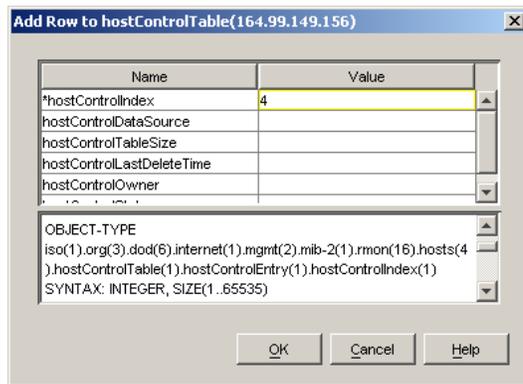
26.3.6 Forming Tables of Scalar Objects

You can make a scalar table by combining the scalar objects from the MIB Browser. A scalar table is a two-column table with the name and value of the scalar object entries. To create a scalar table, you select a group node with scalar child nodes or a group of scalar objects. For example, you add one or more scalar objects such as ipInDelivers and SysUpTime to make a new scalar table labeled ipInDeliversTable.

If you want to view the scalar tables that you create, save the scalar table as a profile. You can load the scalar table profiles when required.

Figure 26-6 shows the Scalar Table window:

Figure 26-6 *Scalar Table*



To combine scalar objects as a scalar table and view the table:

- 1 Create a new scalar table.
- 2 Add to or modify the existing table by adding scalar entries or by removing entries from the table.
- 3 Save the scalar table as a profile.
- 4 Launch the profile.

To create a new scalar table:

- 1 In Novell ConsoleOne, click the target SNMP-manageable node.
- 2 Click *File > Action > SNMP MIB Browser*.
- 3 Right-click a scalar group or a scalar object then click *New*, and then click *Scalar Table*.

To add or remove scalar entries to an existing table:

- 1 Open an existing scalar table.
- 2 Toggle to the MIB Browser window then click *Add to*, and then click *Scalar_table_name*.
Alternatively, click the scalar entry in the MIB Browser window, and from the Scalar Table window, click *Add Node Selected* from Browser Window.

To remove the scalar entry, click the scalar entry in the Scalar Table window, and click *Remove the Node Selected in This Window*.

26.3.7 Graphing SNMP Request Results

You can plot the SNMP request results in a graph that displays the polled data of the MIB objects. Only attributes of ASN.1 type Integer, Counter, Time Ticker, and Unsigned Integer are plotted as current absolute values.

You can plot more than one object in the same graph, add more objects, or remove the MIB objects from the existing graph. If you want to view the graphs that you create, save the graph as a profile. You can then load the graph profiles when required.

To graph SNMP request results of one or more nodes:

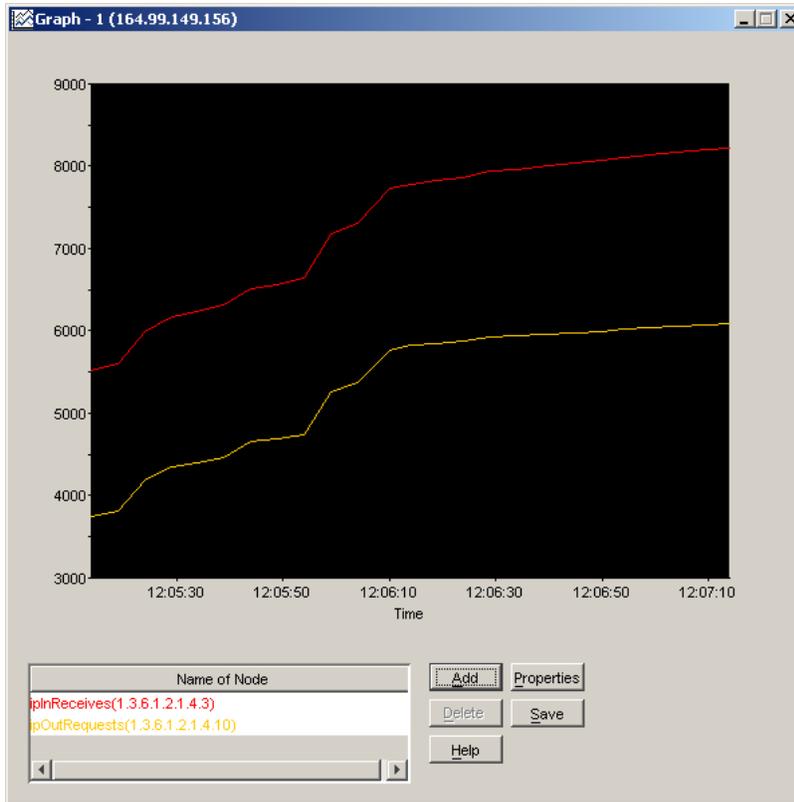
- 1 Click the target SNMP-manageable node from the console.

- 2 Click *File > Action > SNMP MIB Browser*.
- 3 Click the MIB object whose values you want to plot.
- 4 Right-click the object then click *New*, and then click *Graph*.

The MIB Browser plots the graph with the values of the selected object and its leaf object values dynamically in the Graph pane of the window.

Figure 26-7 shows the Graph window:

Figure 26-7 Graph window



To graphically plot the values of more than one object:

- 1 Toggle to the MIB Browser window.
- 2 Click the MIB object you want to plot then click *Add To*, and then click the *Graph*.

You add these objects to any of the active graph windows you want.

Alternatively, you can click the MIB object from the MIB Browser window and then click the *Add* button in the MIB Browser Graph. Remove the objects from the list that you do not want by selecting the node from the list and clicking the *Delete* button.

From the Graph window, you can perform the following operations:

- ◆ Rescale the Y-axis of the graph
- ◆ Set the period to display
- ◆ Set the polling interval and refresh rate of the display

By default, the values plotted in the graph are absolute. If you want to view the rate of change of values per second with respect to `sysUpTime`, you must click the *Rate* option. For example, if you click *ipInPackets* and choose the *Rate* option, you can view the values per second.

26.3.8 Using a Profile for Tables and Graphs

A profile contains information about the properties of the graph, table, or scalar table. You use a profile to specify the information, such as the method of display (table or graph) and polling interval.

You create a profile by saving the properties of the table, scalar table, or graph as a profile. You open the profile to view a table, scalar table, or graph of different SNMP-manageable nodes on the segment with the same properties specified in the profile. You can modify or delete the profile.

To form a profile:

- 1 Save the properties of the display window.
- 2 Open the profile.
- 3 Modify the properties of the profile as required.

To save a profile:

- 1 Click the *Save* button from the Scalar table window, Graph window, or Table window.
- 2 Enter the details of the profile.
Specify the name, description, and properties of the objects.
- 3 Click *OK*.

To open a profile:

- 1 From the MIB Browser window, click the profile you want from the drop-down list.
- 2 Click *Launch This Profile*.

To modify the selected profile:

- 1 Click *View/Edit Profile Contents* for the selected profile in the MIB Browser window.

To delete a selected profile:

- 1 Click *Delete This Profile*.

26.4 Maintaining MIBs

Depending on your need to add MIBs for managing nodes, you must compile the MIBs.

To delete a particular MIB from Novell ZENworks Server Management, remove the appropriate MIB text file from the MIB Server Pool and rerun the MIB Compiler. If the MIB you delete contains traps, you must remove the alarm definitions before you rerun the MIB Compiler.

When you add MIBs, you choose to integrate or exclude the trap information while compiling MIBs. If you disallow trap integration with the MIBs, only the MIB information is stored in the database on successful compilation of the MIBs.

For more information about how to add or remove MIBs, refer to [“Compiling MIBs for SNMP-
Manageable Nodes”](#) on page 1006.

Using the Probe Manageability Tool

27

The Probe Manageability tool helps you find the MIBs implemented on a node or a set of nodes either at the Atlas or the Custom Atlas level.

The following sections provide detail information about how to access and work with the Probe Manageability tool:

- ♦ [Section 27.1, “Invoking the Probe Manageability Tool,” on page 1019](#)
- ♦ [Section 27.2, “Working with the Probe Manageability Tool,” on page 1020](#)

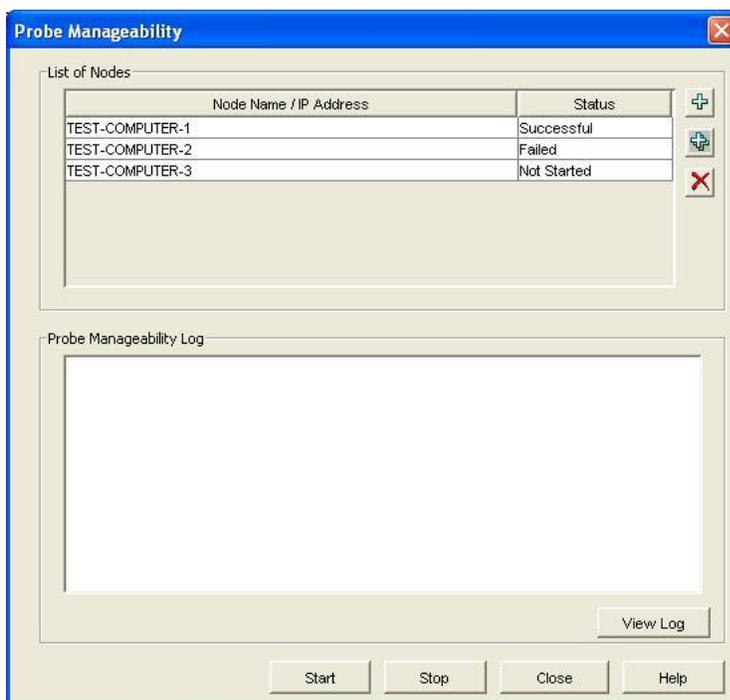
27.1 Invoking the Probe Manageability Tool

You can invoke the Probe Manageability tool in one of the following ways:

- ♦ In ConsoleOne, right-click the Atlas or the Custom Atlas object, click Actions, then click Probe Manageability.
- ♦ In ConsoleOne, select the Atlas or the Custom Atlas object, click the File Menu, click Actions, then click Probe Manageability.

The Probe Manageability window is displayed:

Figure 27-1 Probe Manageability Window



If the Probe Manageability tool is launched from a node level, the List of Nodes pane displays DNS name, IP address, or the Novell eDirectory name of the selected node. It also displays the status of the Probe Manageability operation. The status could be one of the following:

- ♦ **Not Started:** Indicates that the Probe Manageability operation has not yet been performed on the node.
- ♦ **Completed:** Indicates that the Probe Manageability operation has been successfully completed. If Probe Manageability is successful for a single node, then the list of implemented MIBs in it is shown in the Probe Manageability Log pane.
- ♦ **Failed:** Indicates that the Probe Manageability operation has not been successfully completed.

Possible Causes:

(1) SNMP communication failure: The Probe Manageability Log pane displays the SNMP parameters for the node.

(2) Unable to find the node details in the database: The Probe Manageability Log pane displays the node not found in the database message. The complete messages are logged in the file `Novell_ConsoleOne_Directory\1.2\bin\Probe.log`.

This log file contains the list of implemented and unimplemented MIBs and their OIDs for the nodes. It also contains the reason for failure in case probe manageability fails for a node.

27.2 Working with the Probe Manageability Tool

Review the following sections to understand how to work with the Probe Manageability tool:

- ♦ [Section 27.2.1, “Adding a Node to the List,” on page 1020](#)
- ♦ [Section 27.2.2, “Adding Multiple Nodes at a Time to the List,” on page 1020](#)
- ♦ [Section 27.2.3, “Deleting a Node from the List,” on page 1021](#)
- ♦ [Section 27.2.4, “Starting the Probe Manageability Operation,” on page 1021](#)
- ♦ [Section 27.2.5, “Stopping the Probe Manageability Operation,” on page 1021](#)
- ♦ [Section 27.2.6, “Viewing the Probe Manageability Log for a Node,” on page 1021](#)

27.2.1 Adding a Node to the List

- 1 In the Probe Manageability window, click  located in the List of Nodes pane.
- 2 In the Add a Node window, type the eDirectory name, the DNS name or the IP address of the node. You can also click to browse to and select a node.
- 3 Click Add.

The node is added to the list of nodes.

27.2.2 Adding Multiple Nodes at a Time to the List

- 1 In the Probe Manageability window, click  located in the List of Nodes pane
- 2 In the Select Objects window, select multiple nodes that you want to add to the list.
- 3 Click OK.

The selected nodes are added to the list of nodes.

27.2.3 Deleting a Node from the List

- 1 In the Probe Manageability window, select the node you want delete from the list of nodes
- 2 Click .

The selected node is deleted from from the list of nodes.

27.2.4 Starting the Probe Manageability Operation

- 1 In the Probe Manageability window, click Start.

27.2.5 Stopping the Probe Manageability Operation

- 1 In the Probe Manageability window, click Stop.

27.2.6 Viewing the Probe Manageability Log for a Node

- 1 Click View Log.

The log for nodes for which Probe Manageability is completed till that instance is shown in the Log pane. However, to view the information of a specific node, select the node in the List of Nodes pane. At the end of current Probe Manageability session, log for all the nodes is displayed in the Log pane.

The location of log file is

Novell_ConsoleOne_Directory\1.2\reporting\Probe_Manageability\Probe.log:

Figure 27-2 Probe Manageability Log

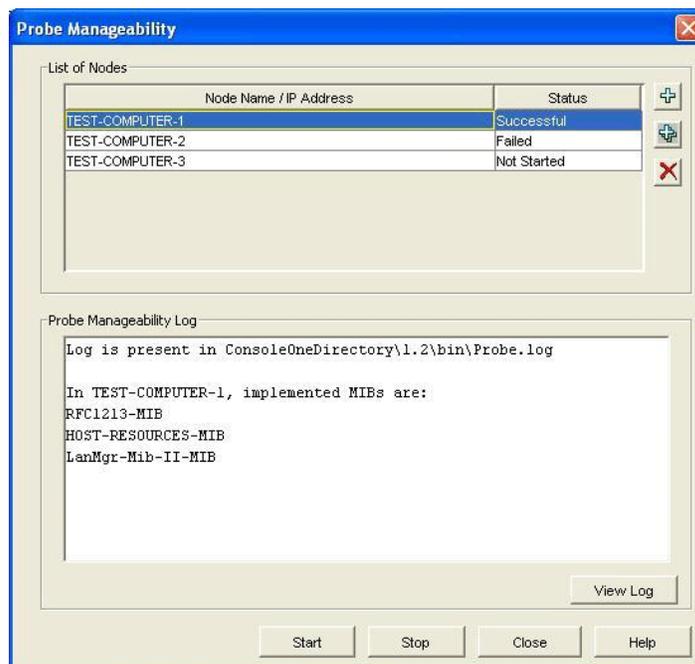
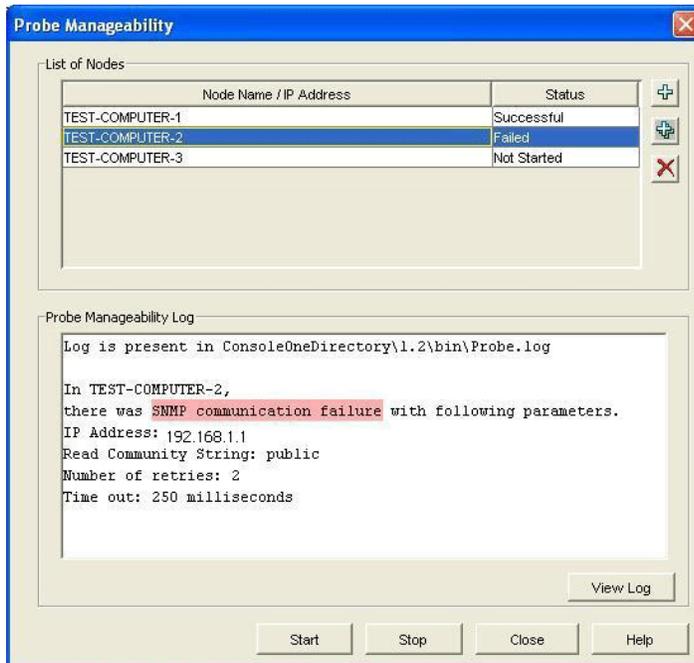


Figure 27-3 Probe Manageability Log with details of SNMP failure.



Novell ZENworks[®] Server Management lets you test the connectivity and availability of a service on a network device. This test checks and measures the response by sending diagnostic packets, and also notifies the console whenever the status of the service changes.

This section provides an overview of the testing facility, lists the services that can be monitored on the nodes, and discusses the test options. See the following sections for more information:

- ♦ [Section 28.1, “Understanding Monitoring Services,” on page 1023](#)
- ♦ [Section 28.2, “Monitoring Services on Target Nodes,” on page 1025](#)

28.1 Understanding Monitoring Services

Using the Monitoring Services facility, you test connectivity of services on one or more critical network devices, such as servers or routers. For example, you can monitor services because you want to be alerted immediately if the connectivity between the console and critical nodes is disrupted.

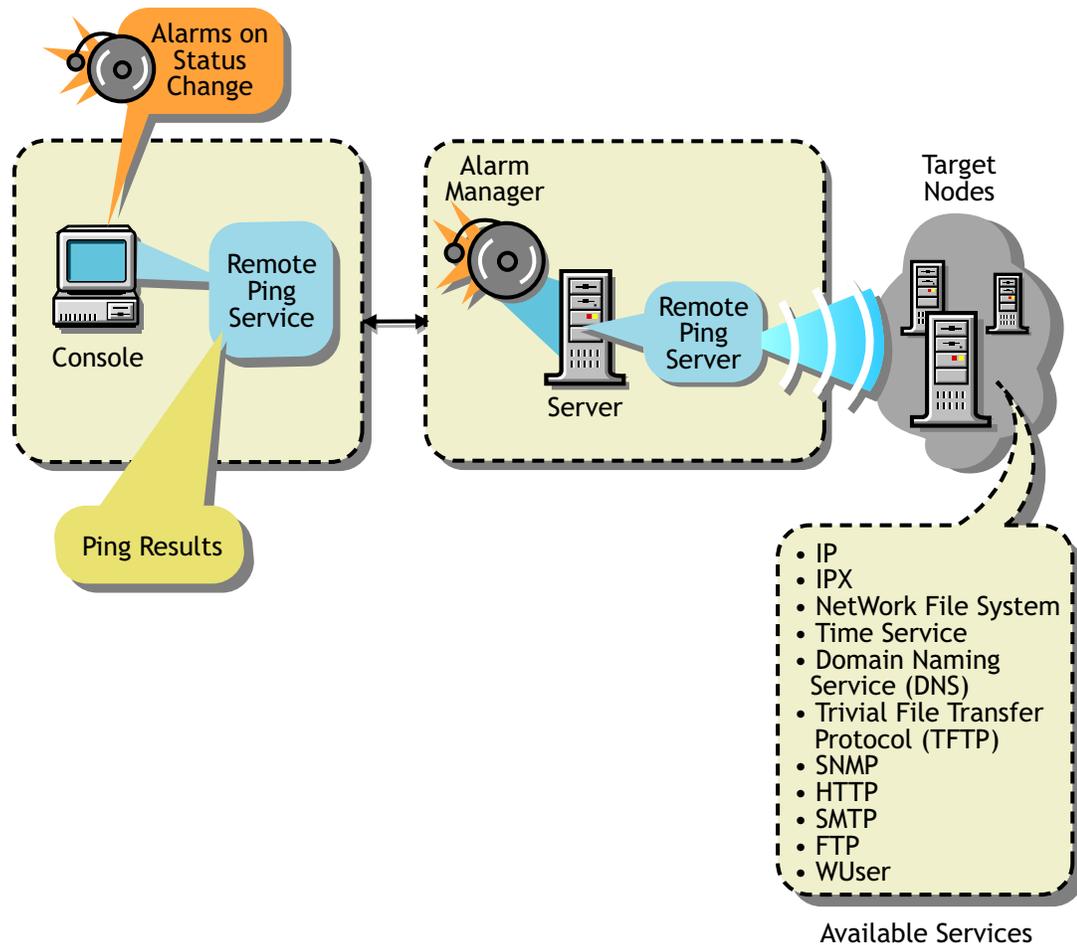
This test facility enables testing of the following services:

- ♦ Domain Name System (DNS)
- ♦ Dynamic Host Configuration Protocol (DHCP)
- ♦ Echo
- ♦ File Transfer Protocol (FTP)
- ♦ Hypertext Transfer Protocol (HTTP)
- ♦ Hypertext Transfer Protocol Secure (HTTPS)
- ♦ Internet Packet Exchange™ (IPX™)
- ♦ Internet Protocol (IP)
- ♦ Network File System (NFS)
- ♦ Network News Transfer Protocol (NNTP)
- ♦ Simple Mail Transfer Protocol (SMTP)
- ♦ Simple Network Management Protocol (SNMP)
- ♦ Time Service
- ♦ Trivial File Transfer Protocol (TFTP)
- ♦ WUser

The test facility uses the Novell ZENworks Server Management server as the remote ping server. When you select the service on the node for testing, the console interacts with the remote ping server on the Novell ZENworks Server Management server and displays the results of the test on the console.

[Figure 28-1](#) shows a graphical representation of Monitoring Services:

Figure 28-1 The Console, Remote Ping Server, and the target nodes



To monitor nodes, you choose the nodes and enable the monitoring session for the duration you require.

From the console, you monitor the services in the following ways:

- ◆ Test connectivity of the services on a node one time only when you suspect a problem with the connectivity.
- ◆ Continuously monitor connectivity of the services on a critical node until you close the test facility.
- ◆ Continuously poll the services of the nodes on the segment (for example, connectivity testing of the services on the target nodes runs uninterrupted until you disable monitoring). If you do not disable monitoring, this test facility continues even after you close the console.

For testing connectivity of services on the target nodes you select, you set the following options:

- ◆ Specify the services on the selected target nodes.
 - If you need to test any TCP-based services, add the service to the existing list of services.
- ◆ Define the test interval between two successive tests.
- ◆ Define the timeout value.

The timeout value determines the time duration that the remote ping server waits to receive the response from the target node.

You can view the status of the connectivity and measure diagnostics, such as round trip delays or number of packets sent and received from the console.

28.1.1 Role-Based Services for Using the Monitoring Services

Role-based services (Role-based Services) defines the task for Monitoring Services as Enable Remote Ping. If this task is assigned to your role, you can use the Monitoring Services facility.

For general information about role-based traffic analysis tasks, creating Role-based Services role objects or specifying tasks that Role-based Services roles can perform, see [Section 21.3, “Role-Based Administration,” on page 830](#).

28.2 Monitoring Services on Target Nodes

This section guides you through the tasks involved in using the Monitoring Services facility.

From the console, you can monitor critical nodes on the network and manage potential connectivity problems before they affect the network. You define the services to test on the selected nodes, then view the test results and other data for each listed target. To perform the testing, complete the following general steps:

1. Define the targets to be monitored.

See [“Defining the Targets for Monitoring Services” on page 1025](#) for information about specifying the services on the target nodes.

2. On a per-node basis or on multiple nodes, change the test interval or timeout value.

These tests use default values for the test interval between two successive tests on the target and to determine the time duration that the remote ping server waits to receive the response from the target node. You can change these values for the test.

See [“Changing the Test Options for a Node” on page 1028](#) for information about editing the test options.

3. View the test results.

The nodes are monitored continuously, at the defined test interval for the node. Depending on the Monitoring Services test that you choose, the corresponding test results are displayed.

See [“Displaying Test Results Data” on page 1027](#) for information about test results data.

28.2.1 Defining the Targets for Monitoring Services

Monitoring Services requires that you specify the targets for the tests. You can choose from the following test options:

- ♦ [“Test the Services on the Target Node One Time Only” on page 1026](#)
- ♦ [“Continuously Monitor the Services on the Target Nodes” on page 1026](#)
- ♦ [“Continuously Poll the Services of the Target Nodes on a Segment Until the Test Is Disabled” on page 1026](#)

NOTE: You can monitor approximately 50 critical services simultaneously on the servers. Monitoring more than 50 services may overload the server memory and result in performance degradation.

Test the Services on the Target Node One Time Only

If you suspect a problem with a node in the network, you can ping the node once for monitoring services. When you select the target node for testing the services and specify the IP or IPX address, this address will determine the service that will be tested at the node. For example, if you enter the IPX address, the default IPX service is tested on the target node.

The results of the test will display the status of the target node and details of the round trip delay in the Ping window.

To test the services on a node once:

- 1 In Novell ConsoleOne, right-click the selected node, then click *Ping*.
- 2 Enter the ping target details.
- 3 Click *OK*.

Continuously Monitor the Services on the Target Nodes

To specify the services for continuous monitoring, add the targets and choose the services on the node and other options. The target node will be added to the list of targets in the Connectivity Test Results window and the test results data will be displayed. Monitoring of services continues until you close this window.

To define the targets for testing services on the node:

- 1 In Novell ConsoleOne, click *Action*, then select *Connectivity Test*.
- 2 Click *Add*.
- 3 Specify the details for the target nodes in the Add Ping Target dialog box.
Refer to [“Adding Services for Monitoring” on page 1028](#) for more information about adding services.
- 4 Click *OK*.

The target node will be added to the list of targets in the Connectivity Test Results window.

Continuously Poll the Services of the Target Nodes on a Segment Until the Test Is Disabled

For polling the services on the nodes of a segment, select the nodes on a segment with the list of services you want to test. Enable the test in the Monitor Tab Services window and view the results of the test in the Polling view.

If you do not disable the test, polling of the services continues after you close the console.

To define the services on the nodes for polling:

- 1 In Novell ConsoleOne, right-click the node of a segment then click *Properties*, and then click the *Monitor Services* tab.

The List of Segment dialog box displays the different addresses of the same node on different segments if the node is connected to more than one segment. Click the node on the segment that you want to add.

- 2 Specify the details for the target nodes in the Monitor Services Tab window.

Refer to [“Adding Services for Monitoring” on page 1028](#) for more information about adding services.

- 3 Click *OK*.

28.2.2 Displaying Test Results Data

After defining the services for testing on the target node, you can view the results from the console.

Depending on the test you choose, the test results are displayed in the corresponding window.

If you choose to test the services on the node one time only, the test results will be displayed in the Ping Status window of the Ping window. This target will not be tested in the Connectivity Test Results window.

If you choose to continuously monitor the services, the test continues until you close the window. You can view the results in the Connectivity Test window.

If you choose to continuously poll the services until you disable the test, you can view the test results in the Polling view.

The following test data is available when you monitor the services on the target nodes:

Ping Target: Name or address (IP or IPX) of the network device for which services are being tested.

Service: Monitored services that are being tested on the target.

Port: Port number that the service uses.

Status of the Target: Up Status means that the service is available on the node and can be reached from the remote ping server. Down Status means that the service is down and cannot be reached from the server.

RoundTrip Delay: Time interval (in milliseconds) between the instant the remote ping server sends the test packet to the target and the instant the response is received from the target.

Packets Sent: Number of packets sent from the remote ping server to the target node.

Packets Received: Number of packets received by the remote ping server from the target node.

Packets Lost: Number and percentage of packets lost during the testing of the target node.

Interval: Displays the test interval value, in seconds. This value determines the time duration between two successive tests on the target.

Timeout: Displays the timeout value, in milliseconds. This value determines the time duration that the remote ping server waits to receive the response from the target node.

To view the Connectivity Test Results window:

- 1 Click *File > Action > Connectivity Test* from the Console.

If you select one or more target nodes from the right pane of the console, the list of nodes that you want to test for connectivity will be shown in the Connectivity Test Results window.

To view the results of the polling:

- 1 From the console, click a segment, then *View > Polling*. The result would also appear at the atlas level. Select any atlas, click *View > Polling*. You can use this option to view the status of all the nodes that you have added using the Monitor Services tab.

NOTE: To delete a target node from the list, from the Polling view, click the target node, then click *Delete*.

28.2.3 Changing the Test Options for a Node

You can modify the test options, such as the test interval and timeout options, that you set earlier on an individual node or on multiple nodes. To modify multiple nodes, click more than one node from the Connectivity Test Results window; the test options apply to all selected target nodes.

To view the Connectivity Test Results window:

- 1 Click the target row from the Connectivity Test Results window, then click the *Edit* button.
- 2 Enter values for the Ping Interval and Timeout.
- 3 Click *OK*.

If you want to roll back to the default setting, click *Apply Defaults*.

28.2.4 Adding Services for Monitoring

Monitoring Services lets you test services on the nodes. If you need to test any TCP-based service that is not listed in the default services list, you add the details of the service when you are adding the targets.

You specify the name of the service in the Add Service dialog box. Ensure that the service name you add is a unique name. Also, you must specify the port number for the service.

You can add the details of the service under the following circumstances:

- ♦ [“Continuously Monitor the Services on the Target Nodes” on page 1026](#)
- ♦ [“Continuously Poll the Services of the Target Nodes on a Segment Until the Test Is Disabled” on page 1026](#)

The services that you add are stored in a file on the server.

Novell ZENworks[®] Server Management provides traffic analysis tools that monitor network traffic, capture data, and collect key statistics of monitored segments nodes, and devices, allowing you to obtain, review, and analyze vital information to effectively troubleshoot and manage your LAN and keep your network operating at peak performance.

This section contains the following topics:

- ◆ [Section 29.1, “Understanding Traffic Analysis,” on page 1029](#)
- ◆ [Section 29.2, “Planning for Segment Monitoring,” on page 1042](#)
- ◆ [Section 29.3, “Preparing to Analyze Network Traffic,” on page 1044](#)
- ◆ [Section 29.4, “Analyzing Network Traffic,” on page 1047](#)
- ◆ [Section 29.5, “Optimizing Traffic Analysis,” on page 1078](#)
- ◆ [Section 29.6, “Understanding the Traffic Analysis Agents,” on page 1089](#)
- ◆ [Section 29.7, “Using the Traffic Analysis Agent for NetWare,” on page 1090](#)
- ◆ [Section 29.8, “Using the Traffic Analysis Agent for Windows,” on page 1104](#)

29.1 Understanding Traffic Analysis

This section contains basic information to help you understand traffic analysis and describes the Novell ZENworks Server Management traffic analysis components.

- ◆ [Section 29.1.1, “Traffic Analysis Components,” on page 1029](#)
- ◆ [Section 29.1.2, “Communication Between Traffic Analysis Components,” on page 1030](#)
- ◆ [Section 29.1.3, “Traffic Analysis Features,” on page 1031](#)
- ◆ [Section 29.1.4, “Traffic Analysis Fundamentals,” on page 1032](#)

29.1.1 Traffic Analysis Components

The Novell ZENworks Server Management traffic analysis components include:

- ◆ [“Management Server” on page 1029](#)
- ◆ [“Management Console” on page 1030](#)
- ◆ [“Monitoring Agent Server” on page 1030](#)

Management Server

The management server comes with the robust and highly scalable Sybase* Adaptive Server Anywhere that stores static information, such as the names and addresses of the nodes and devices in your network. The management server components include the NetExplorer™, management database, Consolidator, and Atlas Manager. NetExplorer discovers the objects in your network and stores them in the management server. The Consolidator takes the information about network objects discovered by NetExplorer and builds the management database. For details about the functionality of NetExplorer, see [Section 23.1, “Understanding Network Discovery,” on page 864](#).

The management database is comprised of the Common Information Model (CIM) schema that is used to establish the topology of the network. The CIM schema extension capabilities provide the ability to organize the information in the database and give this information the shape of a network map. The Atlas Manager obtains information from the management database and displays the network map on Novell ConsoleOne.

Management Console

Novell ConsoleOne[®], the Novell[®] directory-enabled, Java*-based network management and administration tool, is the management console component. Novell ZENworks Server Management snaps in to Novell ConsoleOne and expands Novell ConsoleOne's capabilities by adding menu options, property pages for existing Novell[™] objects, and ways to browse and organize network resources. Novell ConsoleOne provides an intuitive, graphical user interface for Novell ZENworks Server Management traffic analysis. For details about the functionality of Novell ConsoleOne, see [Section 23.3, “Managing the Atlas,” on page 908](#).

Monitoring Agent Server

Before you start analyzing segments or devices on your network, you need to ensure that they are monitored. To enable monitoring, make sure you have installed the network monitoring agent software either on the management server or on an independent server in your network. For more information, see [“Management and Monitoring Services Installation”](#) in the *Novell ZENworks 7 Server Management Installation Guide*. Network monitoring agents gather information or provide services that help you monitor your network.

An agent program using parameters you have provided searches all or part of your network, gathers information you query, and presents it to you when you require it. You can use the information gathered by the agent to analyze the traffic on your network. The agent also warns you of problems, such as duplicate IP addresses, by sending an alert to Novell ConsoleOne to help you solve problems before network performance is impacted. For details about managing alarms, see [Section 24.2, “Managing the Alarm Management System,” on page 924](#).

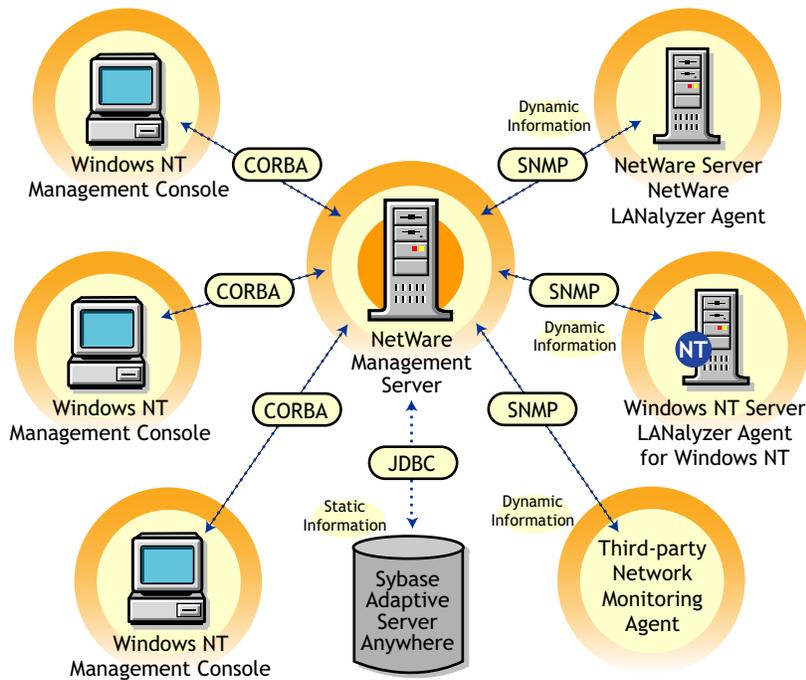
Network monitoring agents observe traffic and capture frames to build a database of network objects and information to help you detect network aberrations. With the network monitoring agent software installed on a server on each of your segments, you can use the traffic analysis tools to help you monitor the traffic on your network, identify the source of network problems, and maintain optimum performance. For details, see [“About Network Monitoring Agents” on page 1032](#). The traffic analysis agents for Novell NetWare[®] and Windows* are part of Novell ZENworks Server Management that you can use to monitor Ethernet, FDDI, or token ring networks.

29.1.2 Communication Between Traffic Analysis Components

Novell ConsoleOne communicates with the management server using common object request broker architecture (CORBA) to procure dynamic and static information about the nodes and devices in your network. When Novell ConsoleOne requests static information from the management server, the management server communicates with the management database using Java Database Connectivity (JDBC), gathers the required static information from the database, and provides it to Novell ConsoleOne. When Novell ConsoleOne requests dynamic information from the management server, the management server communicates with the network monitoring agent using SNMP, gathers the required dynamic information, and provides it to Novell ConsoleOne.

[Figure 29-1](#) illustrates this communication:

Figure 29-1 Communication among Traffic Analysis components



29.1.3 Traffic Analysis Features

The Novell ZENworks Server Management traffic analysis components provide the following features:

- ♦ [“Analyze Traffic Generated by Segments” on page 1031](#)
- ♦ [“Analyze Traffic Generated by Nodes Connected to Segments” on page 1031](#)
- ♦ [“Capture Packets, Decode Captured Packets, and Display Captured Information” on page 1032](#)
- ♦ [“Analyze Traffic Generated by Protocols” on page 1032](#)
- ♦ [“Analyze Traffic Generated by Switches” on page 1032](#)

Analyze Traffic Generated by Segments

You can use the traffic analysis tools to collect current and historical segment statistics that can be displayed in real time, stored for later display, or transferred to a database, spreadsheet, or management reporting system. For details, see [“Analyzing Traffic on Segments” on page 1047](#).

Analyze Traffic Generated by Nodes Connected to Segments

The traffic analysis tools allow you to obtain statistical information about nodes on monitored Ethernet, FDDI, or token ring segments, and determine the top nodes on a segment. You can monitor the status of nodes in your network so that you are alerted when a node becomes inactive. You can also view alarms that are generated when preset threshold parameters are exceeded. Alarms that require immediate attention can be forwarded via e-mail to remote users. For details, see [“Analyzing Traffic on Nodes Connected to a Segment” on page 1055](#).

Capture Packets, Decode Captured Packets, and Display Captured Information

You can use the traffic analysis tools to capture packets between nodes on a monitored segment, and you can quickly define a capture filter based on which you want the packets to be captured. After packets are captured, protocols are decoded and displayed in color-coded summary, decode, and hex panes. The information obtained from the captured packets can be used to examine the traffic on the segment and to analyze it. By providing analysis capabilities and advanced protocol decodes, the traffic analysis tools allow you to identify network aberrations and resolve network performance problems. For details, see [“Capturing Packets” on page 1063](#), [“Protocol Decodes Suite Supported by Novell ZENworks Server Management” on page 1041](#), and [“Displaying Captured Packets” on page 1066](#).

Analyze Traffic Generated by Protocols

You can use the traffic analysis tools to determine the distribution of protocols in the network, transport, and application layer of your network, and obtain statistical information of protocols discovered by the network monitoring agent. For details, see [“Analyzing Traffic Generated by Protocols in Your Network” on page 1073](#).

Analyze Traffic Generated by Switches

You can analyze switch traffic by using the traffic analysis tools to determine port statistics of monitored switches. For details, see [“Analyzing Traffic on Switches” on page 1076](#).

29.1.4 Traffic Analysis Fundamentals

Novell ZENworks Server Management provides tools to let you obtain statistical information about segments, nodes, and devices on your network. You can use this information to analyze and manage the performance of traffic on your network to help you keep the network operating smoothly. Novell ZENworks Server Management also provides tools to capture and decode packets between nodes. You can use the decoded information obtained from captured packets to analyze the traffic between nodes.

To be able to analyze the segments and nodes connected to a segment, you need to ensure that the segment is monitored by a network monitoring agent. You choose the agent based on the type of your network. The Novell ZENworks Server Management traffic analysis tools include the Traffic Analysis Agent for NetWare and Traffic Analysis Agent for Windows, which you can use to monitor segments in your network. NetWare 5.x, the management server for Novell ZENworks Server Management, includes Novell eDirectory, which is leveraged by Novell ConsoleOne, to enable role-based administration.

The following sections provide information that will help you understand the Novell ZENworks Server Management traffic analysis functionality:

- ◆ [“About Network Monitoring Agents” on page 1032](#)
- ◆ [“Role-Based Traffic Analysis Tasks” on page 1040](#)
- ◆ [“Protocol Decodes Suite Supported by Novell ZENworks Server Management” on page 1041](#)

About Network Monitoring Agents

Network monitoring agents provide the functionality to remotely monitor segments and devices on your network using SNMP. The agents collect and store statistical and trend information about nodes

and devices on the network to provide real-time information about the status of your network. From your desktop, the agents let you troubleshoot and optimize Ethernet, FDDI, or token ring segments.

Based on the size and type of your network, you can use RMON, RMON Lite, RMON Plus, RMON2, or Bridge agents to monitor traffic. The following sections provide information to help you understand the functionality of agents:

- ◆ “Functionality of RMON Agents” on page 1033
- ◆ “Functionality of RMON Lite Agents” on page 1034
- ◆ “Functionality of RMON Plus Agents” on page 1035
- ◆ “Functionality of RMON2 Agents” on page 1037
- ◆ “Functionality of Bridge Agents” on page 1038
- ◆ “Viewing the Summarized RMON Information” on page 1039

Functionality of RMON Agents

RMON agents use a standard monitoring specification that allows various nodes and console systems on your network to exchange network data. This data can be used by a network administrator to monitor, analyze, and troubleshoot a group of distributed LANs from a central site. RMON is specified as part of the MIB in [RFC 1757](http://www.isi.edu/in-notes/rfc1757.txt) (<http://www.isi.edu/in-notes/rfc1757.txt>) as an extension of the SNMP.

RMON agents are ideally used for monitoring Ethernet, FDDI, or token ring segments.

RMON agents collect information in the nine RMON groups of monitoring elements in [Table 29-1](#), each providing specific sets of data to meet network monitoring requirements. For details, see [RFC 1757](#) (<http://www.isi.edu/in-notes/rfc1757.txt>).

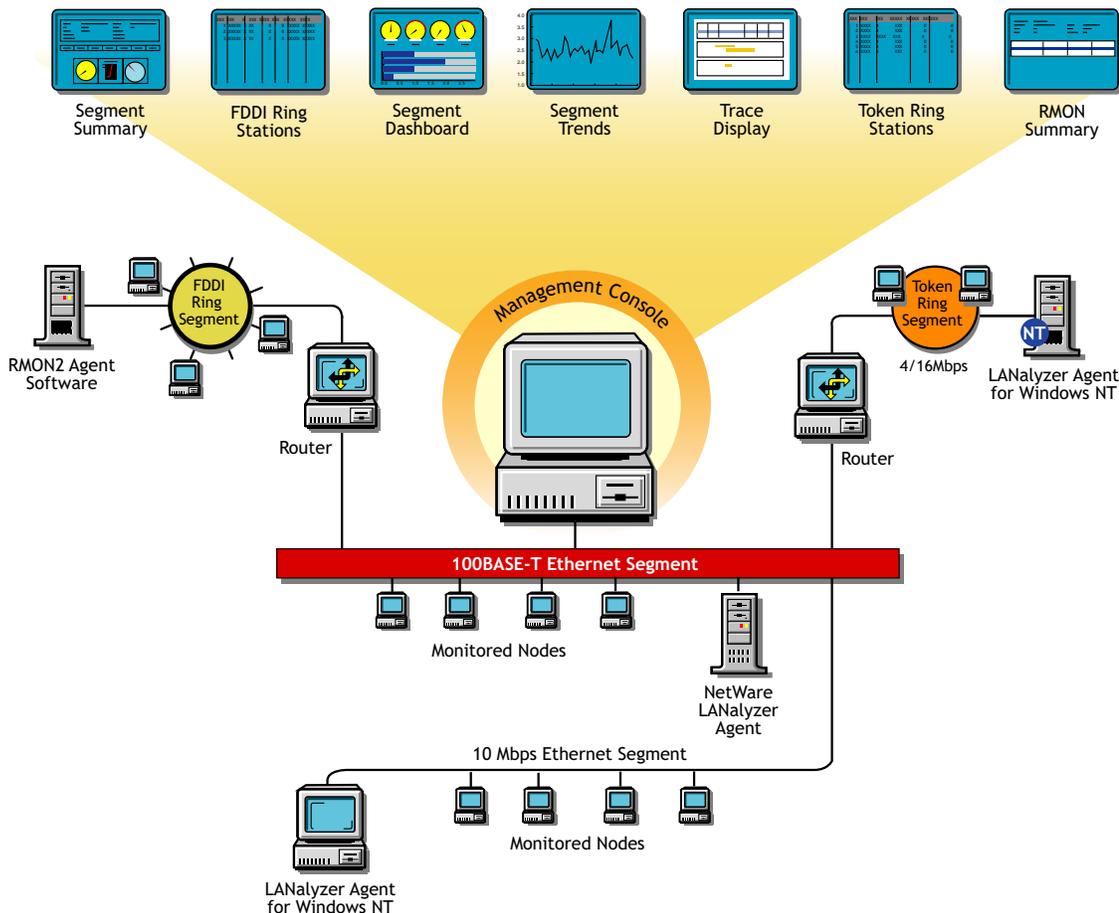
Table 29-1 List of RMON groups of monitoring elements

| RMON Group | Description |
|----------------|---|
| Statistics | Contains statistics measured by the agent for each monitored interface on the device. |
| History | Records periodic statistical samples from a network and stores them for later retrieval. |
| Alarm | Periodically takes statistical samples from variables in the agent and compares them with previously configured thresholds. If the monitored variable crosses a threshold, an event is generated. |
| Host | Contains statistics associated with each host discovered on the network. |
| HostTopN | Prepares tables that describe the hosts that top a list ordered by one of their statistics. |
| Matrix | Stores statistics for conversations between sets of two nodes. As the device detects a new conversation, it creates a new entry in its table. |
| Filters | Allows packets to be matched by a filter. These matched packets form a data stream that may be captured or generate events. |
| Packet Capture | Allows packets to be captured after they flow through a channel. |

| RMON Group | Description |
|------------|---|
| Events | Controls the generation and notification of events from the device. |

Figure 29-2 illustrates the Novell ZENworks Server Management views that you can display when you use an RMON agent to monitor the nodes and devices on your network.

Figure 29-2 Novell ZENworks Server Management views available through an RMON agent



Functionality of RMON Lite Agents

RMON Lite agents are ideally used for monitoring devices not dedicated for network management. For example, RMON Lite agents can be used to monitor a switch in your network.

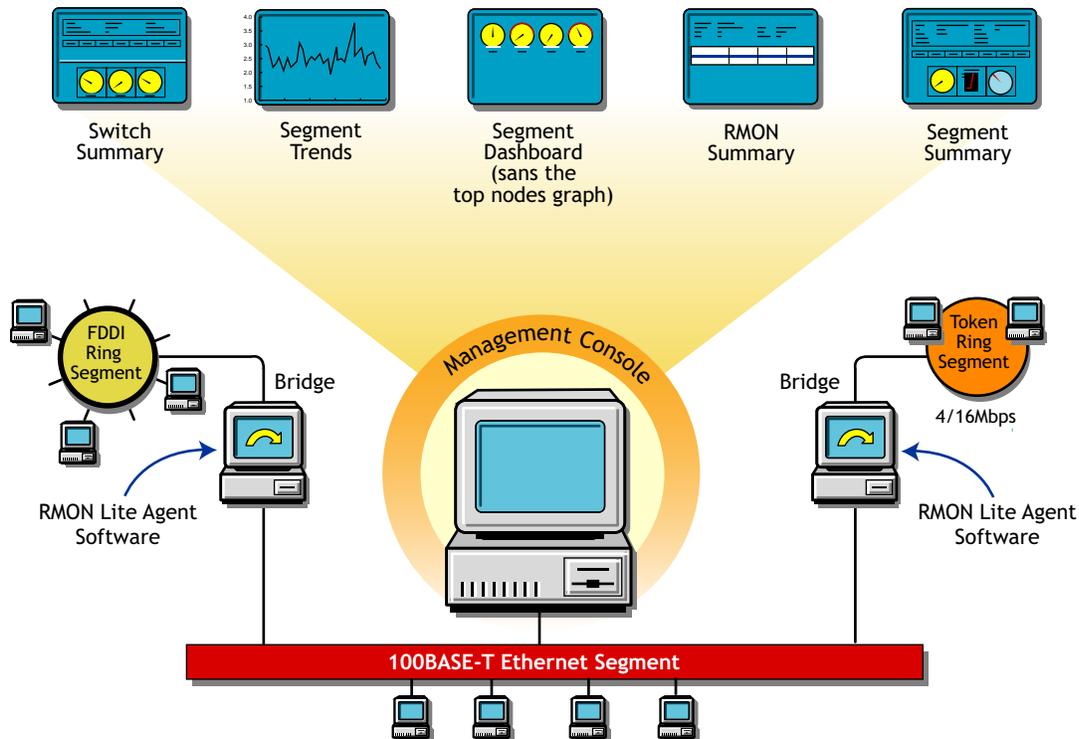
RMON Lite agents support the following four RMON groups:

- ◆ Statistics
- ◆ History
- ◆ Alarm
- ◆ Event

Refer to the table in “Functionality of RMON Agents” on page 1033 for a brief description of each group.

Figure 29-3 illustrates the Novell ZENworks Server Management views that you can display when you use an RMON Lite agent to monitor the nodes and devices on your network.

Figure 29-3 Novell ZENworks Server Management views available through an RMON Lite agent



Functionality of RMON Plus Agents

RMON Plus agents are proprietary agents that extend the functionality of the RMON agent by providing data collected from the RMON groups, explained in “Functionality of RMON Agents” on page 1033, and the groups explained in Table 29-2:

Table 29-2 Functionality of RMON Plus Agents

| RMON Plus Group | Description |
|-----------------|---|
| Buffer | Records the number of octets (excluding framing bits but including frame check sequence [FCS] octets and overhead) in packets which are captured in the buffer. |
| Admin | Collects information specific to the agent, such as the version number. |
| HostMonitor | Monitors a set of nodes for a particular host table and sets traps when a host becomes active or inactive. |
| DuplicateIP | Records and updates a list of packets arriving with duplicate IP addresses. |
| MacToIP | Stores records of the IP addresses associated with a host address for an individual host table. |
| BoardStatus | Records the status of each logical interface of the RMON agent. |

RMON Plus agents are ideally used for monitoring Ethernet, FDDI, or token ring segments. Data from different media types can be collected based on the version of the RMON Plus agent that is used to monitor traffic on your network.

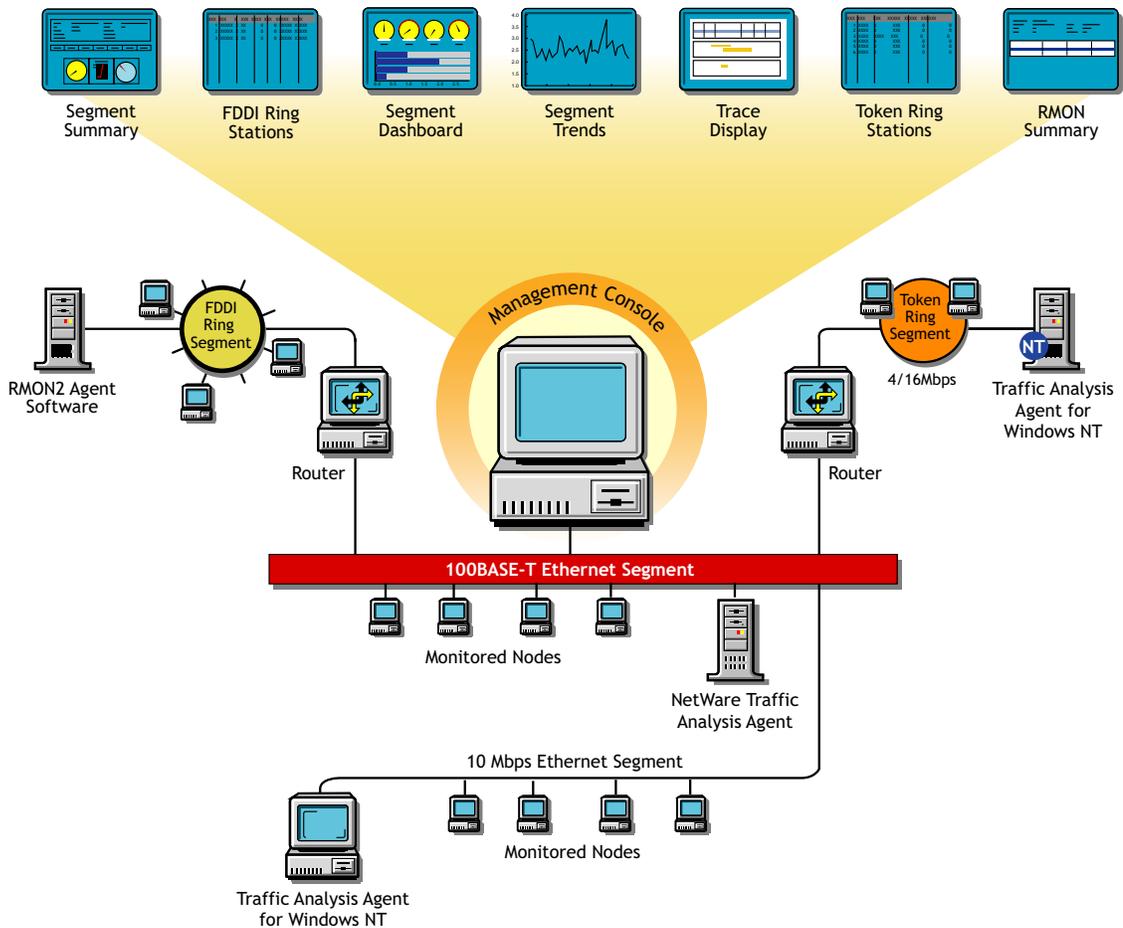
Refer to [Table 29-3](#) to determine the media type support based on the version of the RMON Plus agent:

Table 29-3 List of Media type support based on the version of the RMON Plus agent

| RMON Plus Agent | Media Support |
|---|-------------------------------|
| Traffic Analysis Agent for NetWare 1.1 | Ethernet and token ring |
| Traffic Analysis Agent for NetWare 1.21 or later | Ethernet, FDDI, or token ring |
| Traffic Analysis Agent (version 1.30) for Windows | Ethernet, FDDI, or token ring |

[Figure 29-4](#) illustrates the Novell ZENworks Server Management views that you can display when you use an RMON Plus agent to monitor the nodes and devices on your network.

Figure 29-4 Novell ZENworks Server Management views available through an RMON Plus agent



Functionality of RMON2 Agents

RMON agents can be used to collect data from nodes and devices in the physical and the data link layers and RMON2 agents can be used to collect data from nodes and devices in the network and application layers of your network. RMON2 agents can also determine network usage based on the protocol and application used by the nodes in your network. The RMON2 groups make it possible to view traffic patterns above the data link layer, as shown in [Table 29-4](#). For details, see [RFC 2021 \(http://www.isi.edu/in-notes/rfc2021.txt\)](http://www.isi.edu/in-notes/rfc2021.txt).

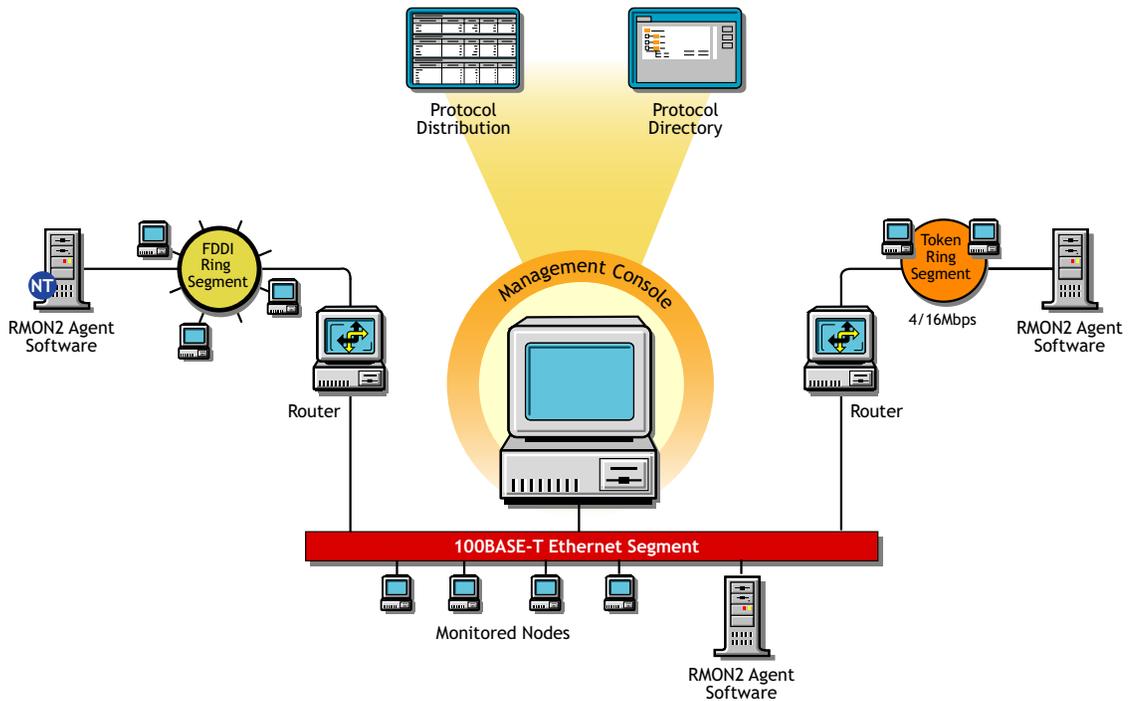
Table 29-4 *Functionality of RMON2 Agents*

| RMON2 Group | Description |
|--------------------------|---|
| Protocol Directory | Provides a table of all identifiable protocols and their descriptions. |
| Protocol Distribution | Provides statistics for each protocol that the agent is configured to track. |
| Address Map | Maps a network layer address to the corresponding Media Access Control (MAC) address. |
| Network-Layer Host | Provides statistics for each host by network layer address. |
| Network-Layer Matrix | Provides statistics for each network conversation between pairs of network layer addresses. |
| Application-Layer Host | Provides statistics on traffic generated by each host for a specified application layer protocol. Traffic broken down by protocols can be recognized by the Protocol Directory group. |
| Application-Layer Matrix | Provides statistics on conversations between pairs of network layer addresses for a specified application layer protocol. Traffic broken down by protocols can be recognized by the Protocol Directory group. |
| User History | Enables the agent to save samples of RMON2 data for any MIB object at specified intervals. |
| Probe Configuration | Provides remote capability for configuring and querying agent parameters such as resets, software updates, IP address changes, and trap destinations. |
| RMON Conformance | Provides information to management software regarding the status of support for the groups. |

IMPORTANT: The Console supports only the Protocol Directory and Protocol Distribution groups.

[Figure 29-5](#) illustrates the Novell ZENworks Server Management views that you can display when you use an RMON2 agent to monitor the nodes and devices on your network.

Figure 29-5 Novell ZENworks Server Management views available through an RMON2 agent



Functionality of Bridge Agents

Bridges are used to connect LAN segments below the network layer. A bridge connects two or more physical networks, forwarding packets between networks based on the information in the data link header.

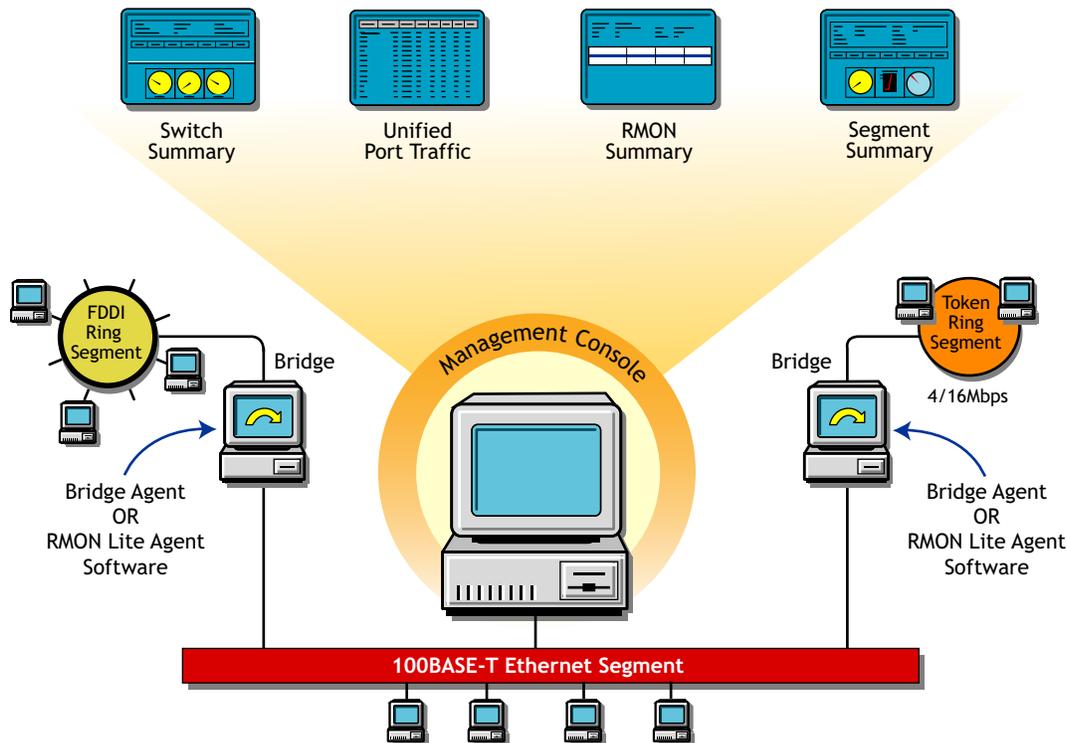
Bridge agents collect information in the five Bridge groups shown in [Table 29-5](#). You can use this information to monitor switched networks. For details, see [RFC 1493](http://www.isi.edu/in-notes/rfc1493.txt) (<http://www.isi.edu/in-notes/rfc1493.txt>).

Table 29-5 Functionality of Bridge Agents

| Group | Description |
|------------------------|---|
| Base | Stores information about objects that are applicable to all types of bridges. |
| Spanning Tree Protocol | Stores information regarding the status of the bridge with respect to the Spanning Tree Protocol. |
| Source Route Bridging | Provides information that describes the status of the device with respect to source route bridging. |
| Transparent Bridging | Provides information that describes the entity's state with respect to transparent bridging. |
| Static | Collects information that describes the entity's state with respect to destination address filtering. |

[Figure 29-6](#) illustrates the Novell ZENworks Server Management views that you can display when you use a Bridge agent to monitor the nodes and devices on your network:

Figure 29-6 Novell ZENworks Server Management views available through a Bridge agent



Viewing the Summarized RMON Information

The RMON Summary view provides brief information about RMON service on a selected node. It displays static information about the RMON agent and details of the resources requested by the user from the agent. The resource requests that are displayed in the RMON Summary view are Packet Capture and Host TopN requests.

To view the summarized RMON information:

- 1 Click *RMON* under *Services* within a node.
- 2 Click *View > RMON Summary*.

Table 29-6 describes the static information displayed in the RMON Summary view:

Table 29-6 Static information displayed in the RMON Summary view

| Statistic | Explanation |
|----------------------|--|
| Agent Name | Name of the RMON agent monitoring the selected segment |
| IP Address | IP address of the node on which the RMON agent is installed |
| IPX™ Address | Internetwork Packet Exchange™ (IPX) address of the node on which the RMON agent is installed |
| Number of Interfaces | Number of logical interfaces for the management server on which the RMON agent is installed |
| Version | Version number of the RMON Plus agent |

| Statistic | Explanation |
|----------------------|---|
| Type of RMON Service | Type of the RMON agent: RMON, RMON Plus, or RMON2 |
| Status of the Agent | Status of the RMON agent |

The RMON Summary view displays the resource information described in [Table 29-7](#):

Table 29-7 Resource information displayed in the RMON Summary view

| Statistic | Explanation |
|---------------|---|
| Resource Name | Type of resource requested: <ul style="list-style-type: none"> ◆ Packet Capture ◆ Host TopN |
| Owner | Owner string corresponding to the control entry of the row |
| Index | Channel, Filter, or Buffer control indexes for the Packet Capture resource and the Control index for the Host TopN resource |

To delete a resource:

- 1 Select a row from the Resource table.
- 2 Click *Delete*.

When you delete a resource, the entry on the agent corresponding to the selected row is deleted.

Role-Based Traffic Analysis Tasks

Novell ZENworks Server Management lets you perform the following traffic monitoring tasks based on your role:

- ◆ Add nodes to be monitored for inactivity.
For details, see [“Monitoring Nodes for Inactivity” on page 1061](#).
- ◆ Add protocols to the protocol directory tree.
For details, see [“Displaying a List of Protocols Used in Your Network” on page 1073](#).
- ◆ Capture packets.
For details, see [“Capturing Packets” on page 1063](#).
- ◆ Disable nodes from being monitored for inactivity.
For details, see [“Monitoring Nodes for Inactivity” on page 1061](#).
- ◆ Delete protocols from the protocol directory tree.
For details, see [“Displaying a List of Protocols Used in Your Network” on page 1073](#).
- ◆ Free agent resources.
For details, see [“Viewing the Summarized RMON Information” on page 1039](#).
- ◆ Set segment alarms.
For details, see [“Configuring Alarm Options from the Set Alarm Dialog Box” on page 1085](#).

- ◆ View conversations.
For details, see [“Viewing Conversations \(Traffic\) Between Nodes”](#) on page 1060.
- ◆ View Traffic Analysis Agents.
For details, see [“Selecting the Preferred RMON Agent”](#) on page 1045.
- ◆ View the protocol directory.
For details, see [“Determining the Distribution of Protocols in a Segment”](#) on page 1075.
- ◆ View the RMON summary.
For details, see [“Viewing the Summarized RMON Information”](#) on page 1039.
- ◆ View segment alarms.
For details, see [“Viewing Alarm Statistics for a Segment”](#) on page 1053.
- ◆ View the segment dashboard.
For details, see [“Determining the Performance of Individual Segments”](#) on page 1049.
- ◆ View segments monitored for inactivity.
For details, see [“Monitoring Nodes for Inactivity”](#) on page 1061.
- ◆ View segment protocol distribution.
For details, see [“Determining the Distribution of Protocols in a Segment”](#) on page 1075.
- ◆ View segment stations.
For details, see [“Listing Statistics for Segments”](#) on page 1048.
- ◆ View the segment summary.
For details, see [“Viewing the Summarized Segment Information”](#) on page 1054.
- ◆ View segment trends.
For details, see [“Analyzing Traffic on Segments”](#) on page 1047.
- ◆ View switch or port traffic.
For details, see [“Viewing Statistics for Ports in a Switch”](#) on page 1076.
- ◆ View the switch summary.
For details, see [“Viewing the Summarized Switch Information”](#) on page 1077.

For more information about role-based services, see [Section 21.3, “Role-Based Administration,”](#) on page 830.

Protocol Decodes Suite Supported by Novell ZENworks Server Management

Novell ZENworks Server Management decodes several protocol suites. Using Novell ZENworks Server Management, you can analyze and troubleshoot problems in the following protocol suites:

- ◆ Novell NetWare Protocol Suite
- ◆ NetWork File System Protocol Suite
- ◆ Systems Network Architecture Protocol Suite
- ◆ AppleTalk* Protocol Suite
- ◆ TCP/IP Protocol Suite

You need to understand these protocols in order to set up packet capture and interpret the results in the Trace Display window. For more information about these protocol suites and decoding support, see [Appendix 31, “Protocol Decodes Suites Supported by Novell ZENworks Server Management,” on page 1135](#)

Novell ZENworks Server Management also enables you to analyze and troubleshoot problems in the following media:

- ◆ Standard Ethernet
- ◆ IEEE 802.3
- ◆ Token Ring
- ◆ FDDI

29.2 Planning for Segment Monitoring

A baseline defines the typical activity of your network. Keeping a baseline document of activity on a segment lets you determine when the activity is atypical. Atypical activity might be caused by a problem or network growth. To create a baseline activity, you should gather statistical information when the network is functioning typically.

The following sections provide information about creating and using a baseline:

- ◆ [Section 29.2.1, “Creating a Baseline of Typical Segment Activity,” on page 1042](#)
- ◆ [Section 29.2.2, “Using the Baseline Document,” on page 1042](#)
- ◆ [Section 29.2.3, “Segment Baseline Document Tips,” on page 1043](#)

29.2.1 Creating a Baseline of Typical Segment Activity

For segment statistics such as bandwidth utilization, you should create a trend graph that plots information over a period of time. Statistics sampling that gathers data over a short period of time can be misleading. If you have added one or more network components, it is useful to create another baseline against which you can compare future activity.

You can export the data you gather in Novell ZENworks Server Management into programs, such as spreadsheets, for further analysis and to maintain records over time.

29.2.2 Using the Baseline Document

You can use the baseline document for the following purposes:

- ◆ [“Using Baseline Documents to Set Alarm Thresholds Appropriately” on page 1042](#)
- ◆ [“Using Baseline Documents to Track Network Growth and Its Effect on Performance” on page 1043](#)
- ◆ [“Using Baseline Documents to Troubleshoot Atypical Segment Activity” on page 1043](#)

Using Baseline Documents to Set Alarm Thresholds Appropriately

Novell ZENworks Server Management lets you set alarm thresholds for statistics on segments monitored by the network monitoring agent software, so that if the threshold is exceeded, you are

notified at Novell ConsoleOne. Setting alarm threshold values for statistics on a segment eliminates the need for you to constantly monitor segments for problems.

Novell ZENworks Server Management provides default values for thresholds of various alarms on Ethernet, FDDI, and token ring segments. Refer to the table in [“Configuring Alarm Options from the Set Alarm Dialog Box” on page 1085](#) for a list of alarm statistics tracked by Novell ZENworks Server Management. By creating a baseline of activity on the segment, you can determine whether the default values are appropriate for segments in your network. For example, after tracking segment utilization, you would set an alarm threshold for bandwidth utilization at about 5% to 10% higher than typical utilization. You are then alerted if utilization is greater than usual for that segment.

IMPORTANT: If you want to use this alarm notification feature, you must enable segment alarms.

Using Baseline Documents to Track Network Growth and Its Effect on Performance

By comparing current network performance against the performance recorded in your baseline document, you can determine how performance is affected by network changes. This comparison also helps you plan for network growth and justify network upgrades and expansion. You can view graphs of real-time trends for various Ethernet, FDDI, and token ring statistics. If an RMON2 agent is installed on a segment, you can also view historical trends for those statistics over hourly, daily, weekly, monthly, and yearly periods. Refer to [“Analyzing Trend Data for a Segment” on page 1050](#) for details about how to view a trend of segment performance. Refer to the table in [“Choosing Options to Display Stations on a Segment” on page 1079](#) for a list of statistics based on which you can display a trend of segment performance.

Using Baseline Documents to Troubleshoot Atypical Segment Activity

By knowing what the typical network activity is, you can recognize atypical activity, which might help you isolate the cause of a problem.

29.2.3 Segment Baseline Document Tips

You should include the following key characteristics in each network baseline document:

- ♦ [“Bandwidth Utilization” on page 1043](#)
- ♦ [“Packets Per Second” on page 1044](#)
- ♦ [“Network Error Rates” on page 1044](#)
- ♦ [“Kilobytes Per Second” on page 1044](#)
- ♦ [“Most Active Servers on the Segment” on page 1044](#)

Bandwidth Utilization

The bandwidth utilization statistic indicates the percentage of network bandwidth used. Bandwidth utilization is likely to be higher at certain times during the day (for example, when users log in to the network in the morning), week, or month. Tracking bandwidth utilization helps you balance traffic loads among network segments, servers, and routers for a more efficient network. This information also helps you determine the effect of network growth on performance. As new workstations and applications are added to a network, bandwidth utilization typically increases.

Packets Per Second

Monitoring the number of packets on the wire provides information about the traffic on the segment. By looking at the change in the packets per second after a user launches a new application, you can calculate what the increase in packets per second will be when all the users you expect to use the application start using it. Packets per second differs from utilization. Utilization is based on the number of kilobytes on the segment per second, but packets can range in size. Therefore, utilization can increase as a result of an increase in the size or number of packets. If the number of packets increases but utilization does not, it is likely that the number of small packets increased but the increase did not affect utilization.

Network Error Rates

By including error rates in your baseline, you can determine when error rates on the network are atypical. This is important because network errors can bring down the network. A higher error rate can result from a hardware problem or network growth. If errors increase but utilization does not, there might be a problem with a component, for example a faulty network board or transceiver.

Kilobytes Per Second

Tracking kilobytes per second lets you determine the throughput of your network. From this information, you can determine the percentage of the total possible bandwidth that is in use. For Ethernet networks, the maximum possible utilization is 10 Mbps. For token ring networks, the maximum possible utilization is 4 or 16 Mbps (depending on the hardware).

Most Active Servers on the Segment

Keeping track of the top three servers on the network helps you distribute the load among them as you add new users and applications. See [“Viewing Statistics of the Top 20 Nodes” on page 1056](#) for details about how to display a list of top nodes on a monitored segment. You should also monitor the number of Request Being Processed packets. A constantly increasing number of these packets indicates a server overload condition. You can monitor these packets by doing a packet capture and decode. See [“Capturing Packets” on page 1063](#) and [“Displaying Captured Packets” on page 1066](#) for details about how to capture and display decoded packets.

With the Segment Trends view, you can view many segment statistics and export that data into another application (such as a spreadsheet) for later analysis. The data is saved as a text file that stores statistical values of the trend you display. To export the trend data to a file, click the Export button in the toolbar of the Segment Trends view. For details, see [“Analyzing Trend Data for a Segment” on page 1050](#).

You can view current utilization for a segment through the Segment Dashboard view. To access this view, select a segment, click View > click Segment Dashboard. For details, see [“Determining the Performance of Individual Segments” on page 1049](#).

29.3 Preparing to Analyze Network Traffic

The Novell ZENworks Server Management software components include the Traffic Analysis Agent for NetWare and Traffic Analysis Agent for Windows. You can install the network monitoring agent on the management server or on an independent NetWare or Windows server. The agent monitors the traffic on the segment it is connected to, gathers information about the nodes and devices on that segment, and makes this information available to the management server, which provides it to Novell ConsoleOne. The agent also sends traps to the management server that are

forwarded to Novell ConsoleOne. The management server and the monitoring agent communicate using SNMP. Novell ZENworks Server Management provides default values for SNMP parameters.

The following sections provide information about specifying a preferred agent for monitoring traffic on the segment and changing the default SNMP settings:

- ♦ [Section 29.3.1, “Selecting the Preferred RMON Agent,” on page 1045](#)
- ♦ [Section 29.3.2, “Setting Up SNMP Parameters,” on page 1046](#)

29.3.1 Selecting the Preferred RMON Agent

If more than one remote monitor (RMON) agent exists on a selected segment, you can choose which agent is to monitor the nodes on the segment from the RMON Agent property page. This page displays a list of servers on which the RMON Agent is installed. The agent installed on the server that you choose from this list becomes the preferred agent. The preferred agent is the primary agent that monitors the segment and sends information about segment activity to Novell ConsoleOne.

To display the RMON Agent property page:

- 1 Select a segment in Novell ConsoleOne.
- 2 Click *File > Properties > the RMON Agent* tab.

[Table 29-8](#) describes the statistics displayed in the RMON Agent property page:

Table 29-8 List of statistics displayed in the RMON Agent property page

| Statistic | Explanation |
|-------------------------|---|
| Preferred | Displays a check mark if the selected server is chosen to be the preferred RMON agent server. |
| Agent Name | Displays a list of all the servers on which the RMON agent is installed. |
| Version | Displays the version of the RMON agent installed on the server. The version is dynamically obtained. If Novell ZENworks Server Management cannot connect to the remote agent, or if a third-party agent is installed on the selected segment, this field is blank. |
| Status | Displays the status of the RMON agent on the segment. |
| MAC Address | Displays the physical Media Access Control (MAC) address of the node. |
| Interface Index | Displays the number of interface indexes in which each interface corresponds to a segment that the node can connect through the network board. |
| Available RMON Services | Displays the list of RMON services available from the selected agent: RMON, RMON Plus, or RMON2. |

To choose an RMON agent as the preferred agent:

- 1 Choose a server or workstation name from the list of names displayed in the property page.

The server and workstation names displayed are those on which the RMON agent is installed.
- 2 Click *Apply*.

29.3.2 Setting Up SNMP Parameters

When you request dynamic information to be displayed in Novell ConsoleOne, it seeks the information from the management server. The management server communicates with the network monitoring agent using SNMP, obtains the required information from the agent, and provides it to Novell ConsoleOne. SNMP communications between the server and the agent are based on default SNMP settings provided by Novell ZENworks Server Management. You can change the default SNMP settings using the SNMP dialog box, which displays in Novell ConsoleOne if an error occurs when the management server is communicating with the monitoring agent.

You can use the SNMP dialog box to specify the community strings and security settings for SNMP communication. You can change the default time-out value for the server to connect with the agent. If the default time-out value is exceeded before the server can communicate with the agent or if the community string of the server does not match that of the agent, the SNMP dialog box displays in Novell ConsoleOne with the current settings. You can use the dialog box to change the current time-out value, the community string, and other SNMP parameters. The changed values are saved in the Novell ZENworks Server Management database and will be applied for all subsequent traffic management sessions.

To change the SNMP settings for all monitoring agents in your network:

- 1 In Novell ConsoleOne, right-click the Novell ZENworks Server Management domain, then click *Global SNMP Parameters*.

To change the SNMP settings for a specific agent:

- 1 In Novell ConsoleOne, right-click the node on which the agent is installed, then click *Properties > SNMP Settings*.

Table 29-9 describes the SNMP parameters displayed in the SNMP Settings property page:

Table 29-9 List of SNMP parameters displayed in the SNMP Settings property page

| Parameter | Explanation |
|------------------|---|
| Community String | Community string of the node requesting dynamic data from the agent |
| Timeout | Maximum duration the server should wait for a response from the agent |
| Retry | Number of times the server should try to connect with the agent |
| Secure Set | Encrypts the packet sent by the management server to the monitoring agent |
| Secure Get | Encrypts the packet sent by the monitoring agent to the management server |

TIP: If the network monitoring agent is running on NetWare 4.x and your network is IPX enabled, use the SNMP dialog box to communicate with the agent using IPX. This will significantly improve the performance of Novell ZENworks Server Management traffic analysis components.

29.4 Analyzing Network Traffic

You can use Novell ZENworks Server Management to monitor your network and collect information such as a summary of real-time statistics to determine the performance of your network, or detailed real-time statistics to determine the performance of segments in your network.

Information about the activity of nodes and segments in your network is presented in views containing tables, dials, and graphs. You can use the information to perform various traffic management tasks such as establishing a baseline on your network to help you identify typical traffic loads and control network problems, and analyze real-time performance to help you balance traffic loads among network segments, servers, and routers. You can also collect node information to help you focus on specific entities that might be the source of problems.

The following sections provide detailed information about how you can use Novell ZENworks Server Management to manage your network monitoring activities:

- ◆ [Section 29.4.1, “Analyzing Traffic on Segments,” on page 1047](#)
- ◆ [Section 29.4.2, “Analyzing Traffic on Nodes Connected to a Segment,” on page 1055](#)
- ◆ [Section 29.4.3, “Capturing Packets,” on page 1063](#)
- ◆ [Section 29.4.4, “Displaying Captured Packets,” on page 1066](#)
- ◆ [Section 29.4.5, “Analyzing Traffic Generated by Protocols in Your Network,” on page 1073](#)
- ◆ [Section 29.4.6, “Analyzing Traffic on Switches,” on page 1076](#)

29.4.1 Analyzing Traffic on Segments

Monitoring the segments on your network helps you keep the network operating cost effectively, consistently, and smoothly. Based on the kind of information you want to obtain, you can choose the agent that will monitor the segments on your network. For details, see [“About Network Monitoring Agents” on page 1032](#). The agent monitoring the segments will collect traffic data and provide real-time or historical information to you when you require it.

Novell ZENworks Server Management provides various views you can use to obtain statistical information about monitored segments. You can choose to view statistical information for all segments in your network or for individual segments. You can view a trend of segment performance and a list of alarms generated on a segment. The Segment Summary view provides a summary of segment performance.

The following sections provide information to help you analyze the performance of segments in your network:

- ◆ [“Listing Statistics for Segments” on page 1048](#)
- ◆ [“Determining the Performance of Individual Segments” on page 1049](#)
- ◆ [“Analyzing Trend Data for a Segment” on page 1050](#)
- ◆ [“Viewing Alarm Statistics for a Segment” on page 1053](#)
- ◆ [“Viewing the Summarized Segment Information” on page 1054](#)

TIP: Servers running the remote monitor (RMON) agent can notify you when nodes you selected for monitoring become inactive. For details, see [“Monitoring Nodes for Inactivity” on page 1061](#). Sometimes the RMON agent server must be taken off the network for maintenance. To prevent the

segment from going unmonitored, you can choose a different RMON agent on the segment. For details, see [“Selecting the Preferred RMON Agent” on page 1045](#).

Listing Statistics for Segments

The List Segments view displays a list of segments and statistical information for each segment on your network. Statistics are displayed in columns of the table in the view. The view displays a list of segments associated with the object or node you selected in Novell ConsoleOne.

See [“Analyzing Traffic on Nodes Connected to a Segment” on page 1055](#) for details about how to use Novell ZENworks Server Management to get information about nodes on individual segments.

To view statistical information of all segments:

- 1 In Novell ConsoleOne, select an Area or a node.
- 2 Click *View > List Segments*.

If you select an Area, the List Segments view displays statistics for all segments found within that Area. If you select a node, statistics for all segments connected to that node will be displayed.

[Table 29-10](#) describes the statistics displayed for each segment. The sampling interval for updating statistics on segments is 15 seconds.

TIP: Statistics of segments are displayed in the List Segments view only if the segments are monitored by a Traffic Analysis Agent for NetWare or Traffic Analysis Agent for Windows.

Table 29-10 *Statistical information of each segment*

| Statistic | Explanation |
|--------------|--|
| Segment Name | Segment name or address. |
| Type | Physical segment type: Ethernet, FDDI, token ring, PPP, and unknown. Unknown indicates the segment whose physical segment type is other than the one listed. |
| Speed (Mbps) | The speed of the segment, as determined by the speed of the network board that attaches the RMON agent to the segment and factors such as the cable type of the segment. The value in this column appears only if you have at least one RMON agent connected to at least one server on your network. |
| Utilization% | Average percentage of the bandwidth currently used by all traffic on the segment. |
| Packets/s | Average number of packets per second currently transmitted on the segment. |
| KBytes/s | Average number of kilobytes per second currently transmitted on the segment. |
| Errors/s | Average number of errors per second currently appearing on the segment. |
| Message | Status of the RMON Agent on the segment. For details, see “Selecting the Preferred RMON Agent” on page 1045 . |

As Novell ZENworks Server Management polls segments, messages in the Messages column vary. These messages display the status of the preferred RMON agent on the segment.

The preferred RMON agent is the node you selected to send information about the segment to Novell ConsoleOne. You can make this selection from the RMON Agent property page. For details, see [“Selecting the Preferred RMON Agent” on page 1045](#).

You can modify the view to show fields; format columns; sort and group items; change the font of text fields; or display grid lines in the table view by selecting the required option from *View > Settings*. For details, see [Chapter 23, “Understanding Network Discovery and Atlas Management,” on page 863](#).

Determining the Performance of Individual Segments

Novell ZENworks Server Management provides real-time statistical information about the monitored segment on your network. This information is displayed in the Segment Dashboard view. The information displayed in this view is useful if you want to troubleshoot a segment.

The Segment Dashboard view displays four gauges that display the real-time statistics for a monitored segment. The lower portion of the view displays a bar graph of the top eight nodes, based on the value selected from the drop-down list. By default, it is based on packets out per second. See [“Viewing Statistics of the Top 20 Nodes” on page 1056](#) for details about how to display a list of the most active nodes on a monitored segment.

You can configure the Segment Dashboard view to display the top eight nodes based on a different statistic. You can also choose to display or disable the top nodes graph. For details, see [“Choosing Options to Display Stations on a Segment” on page 1079](#).

You can set alarm threshold values on segment alarms for packets per second, broadcasts per second, and utilization percentage statistics displayed in the Segment Dashboard view. For details, see [“Defining Alarm Thresholds for Statistics Displayed in the Segment Dashboard View” on page 1050](#).

To view statistical information of an individual segment:

- 1 In Novell ConsoleOne, select a segment.
- 2 Click *View > Segment Dashboard*.

The Segment Dashboard view displays four gauges that display real-time statistics for a monitored segment. The peak value is indicated by a line on each bar in the graph. [Table 29-11](#) describes the statistics displayed in the Segment Dashboard view.

Table 29-11 Statistical information displayed in the Segment Dashboard view

| Statistic | Explanation |
|--------------|--|
| Packets/s | Number of packets per second currently transmitted on the segment |
| Utilization% | Percentage of maximum network capacity currently consumed by packet traffic on the segment |
| Error/s | Number of error packets per second currently transmitted on the segment |
| Broadcasts/s | Number of broadcast packets per second currently transmitted on the segment (a broadcast packet is sent to all addresses on the segment) |

Statistics are updated every five seconds. The numeric value of each statistic is displayed in the gauge.

Defining Alarm Thresholds for Statistics Displayed in the Segment Dashboard View

To set alarm threshold values for statistics displayed in the Segment Dashboard view:

- 1 Click the black ring outlining the gauge.
- 2 Drag the ring to increase or decrease the default values.
- 3 Stop at the value you want to set as the threshold value for the statistic.

As you drag the ring, the color of the ring changes to red.

The color of the ring is displayed in red up to the selected threshold value.

If the statistic on the monitored segment exceeds the threshold value, the RMON agent sends a trap to the management server, which forwards it to Novell ConsoleOne and an alarm is generated.

Viewing the Graph of the Top Nodes on a Monitored Segment

The lower portion of the Segment Dashboard view displays a bar graph of the top eight nodes on a monitored segment. The default statistic on which the graph is based is packets out per second. You can change the statistic on which the graph is based. For details, see [“Choosing the Statistic Based on Which Top Nodes Graph Is Displayed” on page 1083](#). You can also choose to display or disable the top nodes graph. For details, see [“Choosing Options to Display the Top Nodes Graph” on page 1083](#).

Statistics for the graph are updated every five seconds. Every 60 seconds, the graph is re-sorted and the new top nodes are displayed. At this point, new nodes might be added and existing nodes might be discarded from the list.

Analyzing Trend Data for a Segment

Novell ZENworks Server Management allows you to determine trends of traffic patterns on the monitored segment. You can view the trend of segment performance from the Segment Trends view. You can use trend information to create a baseline of typical activity on segments. Having a baseline helps you set appropriate thresholds for segment alarms and plan maintenance activities and backups. Additionally, if problems occur on the segment, you can compare the typical traffic level against the atypical traffic level to help you discover the cause of the problem. For details, see [“Creating a Baseline of Typical Segment Activity” on page 1042](#).

The following topics will help you analyze trend data:

- ◆ [“Understanding the Trend Display” on page 1050](#)
- ◆ [“Viewing Trend Statistics” on page 1051](#)

Understanding the Trend Display

Segment trend data is displayed depending on the type and settings of the RMON agent monitoring the selected segment.

- ◆ If RMON Plus is the segment's preferred RMON agent, you can view current trends gathered every 30 seconds over the last hour and historical trends displayed over hourly, daily, weekly, monthly, or yearly periods.

IMPORTANT: If an RMON agent is installed on more than one node on a segment, the node you select in the RMON Agent property page as the node to send information about the segment to Novell ConsoleOne is the preferred RMON agent server. For more details, see [“Selecting the Preferred RMON Agent” on page 1045](#).

- ◆ If RMON Plus is not selected as the preferred RMON agent for the segment, you can view only the current trends for the selected segment. Current trends are gathered every 30 seconds over the last hour. Select an RMON Plus agent as the preferred RMON agent for the segment to be able to view historical trends.
- ◆ If the preferred RMON agent is Traffic Analysis Agent for NetWare version earlier than 1.30, you can view current trends gathered over the past hour and trends for the past day.
- ◆ Real-time trends will not be displayed if memory usage is excessive or if configuration settings in the RMON agent are unacceptable.
- ◆ If the RMON agent is down or is experiencing problems, the trend for a monitored segment will be displayed as a broken graph.
- ◆ If the preferred RMON agent is a Novell Traffic Analysis Agent (version 1.30 or greater) or a third-party agent that implements the token ring Extensions to the Remote Network Monitoring MIB (RFC 1513), the segment bandwidth utilization graph displays slightly lower values than the actual utilization in the trend for the token ring segment view. This is because the MAC layer statistics are not taken into consideration for the utilization calculation.

Viewing Trend Statistics

To view the trend statistics for a segment:

- 1 In Novell ConsoleOne, select a segment.
- 2 Click *View > Segment Trends*.

Trend graphs are displayed for Ethernet, FDDI, and token ring segments. The default statistics, based on which graphs are displayed for the three types of segments, are shown in [Table 29-12](#):

Table 29-12 *Segment Type Default Statistics*

| Segment Type | Default Statistic |
|--------------|--|
| Ethernet | Total packets, good packets, and error packets |
| FDDI | Total packets |
| Token ring | Total packets |

The toolbar options let you change the time span of the trend you view, select statistics based on which you want the graph to be displayed, and export data to a file.

[Table 29-13](#) describes the toolbar options in detail:

Table 29-13 *Toolbar Options of Segment Trends*

| Option | | Explanation |
|------------------------|---|--|
| <i>Profile</i> |  | <p>Displays the Profile dialog box, from which you can select a default profile. The default profile displays a trend with statistical information for total packets, good packets, and error packets on the monitored segment.</p> <p>If you choose not to use the profiles listed in the <i>Select Profile</i> list, you can select the required statistics from the <i>Select Statistics</i> list. You can save the selected statistics if you want to display the trend of a different segment based on the statistics you selected. The default profile will be enabled the next time you launch the Segment Trends view.</p> |
| <i>Legend</i> |  | <p>Shows what each color in the graph represents. The Legend can be resized.</p> |
| <i>Stack</i> |  | <p>Stacks the trends in a single graph representing all selected statistics, on a single vertical axis.</p> |
| <i>Unstack</i> |  | <p>Un-stacks the trends and displays the graph as a separate strip for each statistic.</p> |
| <i>Horizontal Grid</i> |  | <p>Displays horizontal grid lines in the graph area of the Segment Trend view.</p> |
| <i>Vertical Grid</i> |  | <p>Displays vertical grid lines in the graph area of the Segment Trends view.</p> |
| <i>Scale To Fit</i> |  | <p>Maximizes or minimizes the graph to fit the trend entirely in the graph area of the view.</p> |
| <i>Export</i> |  | <p>Copies the information in the Segment Trends view to a file. The file stores the statistical values displayed by the trend. You can save the data for later analysis.</p> |

| Option | Explanation |
|---|---|
| <i>Time Scale</i> drop-down list  | <ul style="list-style-type: none"> ◆ Real Time: Displays a current trend graph. The default sampling time for this graph is once every minute. This graph updates in real time. ◆ One Hour: Displays a historical graph of the selected trend with a time span of one hour. ◆ One Day: Displays a historical graph of the selected trend with a time span of one day. ◆ One Week: Displays a historical graph of the selected trend with a time span of one week. ◆ One Month: Displays a historical graph of the selected trend with a time span of one month. ◆ One Year: Displays a historical graph of the selected trend with a time span of one year. <p>Historical trends such as hourly, daily, weekly, monthly, and yearly trends are available only when Traffic Analysis Agent for NetWare version 1.1 or later is installed on the segment's preferred traffic analysis agent server.</p> |

The File menu of the Segment Trends view can be used to print the statistical information of the current trend or to export the statistical information of a trend to a file and store the data in text format. You can later import the file into a spreadsheet for analysis.

You can view earlier or ensuing trends and change the size of the graph by using the options available in the graph area of the Segment Trends view, as shown in [Table 29-14](#):

Table 29-14 List of options available in the graph area of the Segment Trends view

| Option | Description |
|---|---|
| <i>Scale Up</i>  | Increments the Y-axis of the graph by half the current size with each click. |
| <i>Scale Down</i>  | Decrements the Y-axis of the graph by half the current size with each click. |
| <i>Previous</i>  | Displays the preceding graph based on the profile or statistics chosen. Enabled only when historical trends are displayed. |
| <i>Next</i>  | Displays the subsequent graph. Enabled only when historical trends are displayed. |

Viewing Alarm Statistics for a Segment

Novell ZENworks Server Management tracks alarm statistics for segments. Alarms are generated when threshold values for statistics on a segment are exceeded. You can view a list of all the alarms for the monitored segment in the Segment Alarms property page.

To view alarm statistics for a segment:

- 1 In Novell ConsoleOne, select a segment.
- 2 Click *File > Properties > Segment Alarms* tab.

Novell ZENworks Server Management provides default threshold values for various segment alarms. You can enable or disable the default values for a monitored segment. If you choose not to use the default values, you can set the threshold value using the Set Alarm dialog box. See [“Configuring Alarm Options from the Set Alarm Dialog Box” on page 1085](#) for details about how to set segment alarms.

If a segment does not have an RMON agent connected to it, an error message is displayed.

Viewing the Summarized Segment Information

The Segment Summary view provides brief information about a monitored segment in your network. It displays static information about the monitored segment, whether the segment is monitored or not, and information about the alarms generated on the segment. At a glance, you can determine the utilization of network capacity by nodes on the monitored segment, view a trend based on packets transmitted by nodes on the segment, and see the distribution of protocols on the segment.

To view the summarized segment information:

- 1 In Novell ConsoleOne, select a segment.
- 2 Click *View > Segment Summary*.

[Table 29-15](#) describes the static information displayed in the Segment Summary view.

Table 29-15 *Static information displayed in the Segment Summary view*

| Statistic | Explanation |
|---------------------|---|
| Name | Name of the segment |
| Type | Media type of the segment: Ethernet, FDDI, or token ring |
| IP Address | IP addresses of the segment |
| IPX Address | IPX address of the segment |
| Primary Agent | Name of the preferred agent monitoring the nodes and traffic on the segment |
| Agent Status | Status of the preferred agent monitoring the nodes and traffic on the segment |
| Nodes | Number of nodes on the segment |
| IP Nodes | Number of nodes on the segment that have an IP address |
| IPX Nodes | Number of nodes on the segment that have an IPX address |
| Servers | Number of NetWare servers on the segments |
| Workstations/Others | Number of nodes on the selected segment that are not NetWare servers |
| Network Probes | Number of monitoring agents on the selected segment |
| Switches | Number of switches on the segment |

| Statistic | Explanation |
|-----------|--|
| Routers | Number of routers used to connect nodes and devices on the segment |
| Hubs | Number of hubs on the segment |

The Segment Summary view displays information about alarms generated on a monitored segment, as described in [Table 29-16](#):

Table 29-16 *List of alarms generated on a monitored segment*

| Statistic | Explanation |
|---------------|--|
| Severity | Severity level attributed to the trap. |
| From | Network address of the device that sent the alarm to the alarm management system. |
| Summary | Summary of the event, often including the name or address of the object affected by the alarm. |
| Owner | Segment or device affected by the alarm. |
| Received Time | Date and time when the alarm management system received the alarm. |
| Type | Generic description of the alarm, for example, Volume out of disk space. |
| Category | Displays the category of the alarm based on the MIB that defines the trap-type objects. The category is directly related to the MIBs included in the management server MIB pool. For example, the category for NetWare servers is based on the NetWare Server Alarm MIB. |

The Segment Summary view displays dynamic information about a monitored segment, as described in [Table 29-17](#):

Table 29-17 *Dynamic information displayed in the Segment Summary view*

| Statistic | Explanation |
|-----------------------|---|
| Utilization% | Displays a dial representing the real-time values of the network capacity consumed by packet traffic on the segment. |
| Packets | Displays the trend based on packets transmitted on the segment. Displays real-time trends for segments monitored by RMON agents and daily trends for segments monitored by RMON Plus agents. |
| Protocol Distribution | Displays a pie chart representing the distribution of application layer protocols for which the agent monitoring the segment can collect data. Each slice represents a protocol suite. Click a slice to view the names of protocols. Enabled if the agent monitoring the selected segment is an RMON2 agent. |

29.4.2 Analyzing Traffic on Nodes Connected to a Segment

Novell ZENworks Server Management provides various views you can use to obtain information about nodes connected to the monitored segments in your network.

The following sections provide information that will help you monitor the performance of nodes connected to the segments in your network:

- ◆ [“Viewing Statistics of the Top 20 Nodes” on page 1056](#)
- ◆ [“Viewing Statistics of Nodes on an FDDI Segment” on page 1057](#)
- ◆ [“Viewing Statistics of Nodes on a Token Ring Segment” on page 1058](#)
- ◆ [“Viewing Conversations \(Traffic\) Between Nodes” on page 1060](#)
- ◆ [“Monitoring Nodes for Inactivity” on page 1061](#)

Viewing Statistics of the Top 20 Nodes

You can use Novell ZENworks Server Management to determine the statistics of the most active nodes on a segment for a wide range of performance statistics. This is useful if you want to discover which node is generating the most traffic based on a particular statistic. For example, you can find the heaviest source of broadcast traffic.

The Stations view displays a list of all nodes on a monitored segment. You can use this view to determine the top 20 nodes on a monitored segment. The view lists the top 20 stations sorted by packets out per second. You can choose a different statistic based on which you want the top 20 nodes to display. For details, see [“Choosing a Statistic Based on Which Top 20 Nodes Are Displayed” on page 1079](#). If there are fewer than 20 top nodes, only the available number of top nodes are listed.

To view the statistics of the top 20 nodes on a segment:

- 1 In Novell ConsoleOne, select a segment.
- 2 Click *View > Stations*.
- 3 From the Stations view, click *View > Show Top N Stations*.

The Stations view displays columns that provide statistical information for each station. [Table 29-18](#) describes the statistics displayed in the Stations view:

Table 29-18 *Statistics displayed in the Stations view*

| Statistic | Explanation |
|---------------|---|
| MAC Address | Physical Media Access Control (MAC) address of a node |
| Node | Name of the node (or address, if the name is not in the database) |
| Util. % | Percentage of maximum network capacity consumed by packets sent by a node |
| Packets/s In | Packets per second received by a node |
| Packets/s Out | Packets per second transmitted by a node |
| Bytes/s In | Bytes per second received by a node |
| Bytes/s Out | Bytes per second transmitted by a node |
| Errors/s | Errors per second transmitted by a node |
| Broadcasts/s | Broadcast packets per second transmitted by a node |

| Statistic | Explanation |
|----------------|---|
| Multicasts/s | Multicast packets per second transmitted by a node (packets transmitted to a specific group of nodes) |
| Protocols | Types of protocols used by a node |
| First Transmit | Date and time a node first transmitted since the traffic analysis agent was started |
| Last Transmit | Date and time a node last transmitted since the traffic analysis agent was started |

Stations statistics are updated periodically. Every 60 seconds, the table is resorted and new top nodes are displayed. At this point, new nodes might be added and existing nodes might be discarded from the list.

Viewing Statistics of Nodes on an FDDI Segment

Novell ZENworks Server Management lets you display data for nodes on monitored FDDI ring segments to help troubleshoot problems.

The FDDI Ring Stations view displays statistics for individual nodes on the monitored FDDI ring segment. The view lists the nodes on the segment and shows the order of each node on the ring and which node is the active monitor.

To view the statistics of nodes on an FDDI ring segment:

- 1 In Novell ConsoleOne, select an FDDI ring segment.
- 2 Click *View > FDDI Stations*.

The statistics shown for each node are cumulative since the Traffic Analysis Agent for NetWare was last started and are updated every ten seconds as described in [Table 29-19](#):

Table 29-19 *Statistics of nodes on an FDDI ring segment*

| Statistic | Explanation |
|---------------------|--|
| Order | Relative position of the node on the FDDI ring from the traffic analysis agent. |
| Name | Name of the node or, if the name is not in the database, the physical (MAC) address of the node. |
| MAC Address | Physical (MAC) address of the node. |
| Status | Status of the node: <ul style="list-style-type: none"> ◆ On—The node is actively participating in a ring poll. ◆ Off—The node is not participating in a ring poll. |
| Duration | Time elapsed since the node was On or Off. |
| UpStream Neighbor | MAC address of the node upstream to this station on the logical ring. |
| DownStream Neighbor | MAC address of the node downstream to this station on the logical ring. |
| Last Entered Time | Date and time the node last entered the ring. |

| Statistic | Explanation |
|-------------------|--|
| Last Exit Time | Date and time the node last exited the ring. |
| SMT Request Type | The SMT request to which the node is responding. Indicates if the node was able to successfully respond to the request. In case of a failure, the response code indicates the reason. |
| SMT Response Type | The SMT response generated by the node on receiving an SMT request. If the node was unable to respond, the response code indicates the reason. |
| Request Denied | The cumulative total of request denied responses generated by the node. A request denied frame is generated when the responding node does not support the SMT version number of the requesting node, when a set fails, or when a request for synchronous bandwidth allocation by a node cannot be honored. |
| In CRC Error | Total number of cyclic redundancy check (CRC) line errors reported by this node. |
| Out CRC Error | Total number of CRC errors reported by the nearest active downstream neighbor of this station and detected by the probe. |
| Lost Frames | Total number of lost frame errors received on the network. A lost frame error indicates that the end delimiter of a frame was lost in the network. |
| In Beacons | Total number of beacon frames detected by the probe that named this station as its upstream neighbor. |
| Out Beacons | Total number of beacon frames sent by this station and detected by the probe. |
| Insertions | Number of times the probe detected this station inserting onto the ring. |

Viewing Statistics of Nodes on a Token Ring Segment

The Token Ring Stations view displays statistics for individual nodes on the monitored token ring segment. The view lists the nodes on the segment and shows the order of each node on the ring and which node is the active monitor.

To view the statistics of nodes on a token ring segment:

- 1 In Novell ConsoleOne, select a token ring segment.
- 2 Click *View > Token Ring Stations*.

The view displays statistical information as described in [Table 29-20](#). Statistics are cumulative since the RMON agent was started and are updated every ten seconds.

Table 29-20 *Statistics of nodes on a Token ring segment*

| Statistic | Explanation |
|-------------|--|
| Order | Relative position of the node on the token ring from the RMON agent. |
| Name | Name of the node or, if the name is not in the database, the physical (MAC) address of the node. |
| MAC Address | Physical (MAC) address of the node. |

| Statistic | Explanation |
|---------------------|---|
| Status | Status of the node: <ul style="list-style-type: none"> ◆ On—The node is on the ring. ◆ Off—The node is off the ring. ◆ On (Monitor)—The node is on the ring and is the active monitor. |
| Duration | How long this node has been on or off. |
| Last Entered Time | Date and time the node last entered the ring. |
| Last Exit Time | Date and time the node last exited the ring. |
| Duplicate Address | Total number of duplicate address errors reported, generated when this node detects other nodes using its own address. |
| Soft Errors | Number of soft errors in packets transmitted by this node. |
| Inline Errors | The total number of line errors reported by this station in error reporting packets to the ring error monitor and detected by the probe. |
| Outline Errors | The total number of line errors reported in error reporting packets sent by the nearest active downstream neighbor of this station and detected by the probe. |
| Internal Errors | Number of internal errors this node has reported. Internal errors generally indicate a recoverable failure of a network adapter board. |
| In Burst Errors | The total number of burst errors reported to the Ring error monitor and detected by the probe. |
| Out Burst Errors | The total number of burst errors reported in error reporting packets sent by the nearest active downstream neighbor of this station and detected by the probe. |
| AC Errors | Number of times this node could not interpret the Address Recognized Indicator (ARI) and the Frame Copied Indicator (FCI) during the ring process. |
| Abort Errors | Number of times a node transmitted an abort sequence. Abort sequences are usually transmitted when a node detects an error in frames it is currently transmitting. |
| Lost Frame Errors | Number of times a node transmitted a frame but failed to receive it back in its entirety. |
| Congestion Errors | Number of times the node detected a frame addressed to its specific address but could not copy it (generally due to insufficient buffers). |
| Frame Copied Errors | Number of times a node detected a frame addressed to its specific address with either or both the ARI and FCI bits set to 1. (Indicates that another node is using its address.) |
| Frequency Errors | Number of times a node's internal clock differed from the ring clock. |
| Token Errors | Number of token errors. These occur when the token gets corrupted or when the Active Monitor does not see a new frame transmitted in the required amount of time. Only the Active Monitor can report this error. |
| In Beacon Errors | The total number of beacon frames sent by this station and detected by the probe. |
| Out Beacon Errors | Total number of beacon frames sent by this station and detected by the probe. |
| Insertions | Number of times the probe detected this station inserting onto the ring. |

| Statistic | Explanation |
|-----------|---|
| Last NAUN | The station that was last named by the probe as the next active upstream neighbor (NAUN). |

Viewing Conversations (Traffic) Between Nodes

Novell ZENworks Server Management provides real-time data about all the network traffic between a selected node and one or more other nodes on a segment. This data can be viewed from the Conversations view. You can use the data displayed in this view to determine specific information about node communication. For example, it can show which nodes communicate with a router or server, determine the load on a server, or examine the traffic flowing to or from a node that is reporting difficulties.

To view conversations between nodes:

- 1 In Novell ConsoleOne, select a node.
- 2 Click *View > Conversations*.

If the selected node is connected to more than one segment, the Select Segment dialog box displays.

- 2a Select the segment where the node you want to examine traffic is connected then click *View*, and then click *Conversations*.

The Conversations view lists the percentage of traffic that each destination node contributes to the load on the source node. However, due to sample skewing (samples not taking place at the same time) and rounding up of statistics, the numbers in the columns do not always add up to 100%.

The statistics displayed in the Conversations view are updated every 5 seconds. [Table 29-21](#) describes the statistics displayed in the Conversations view:

Table 29-21 Statistics displayed in the Conversations view

| Statistic | Explanation |
|-------------|--|
| Node | Name of the destination nodes with which the source node is communicating |
| % Pkt Load | Percentage of the packet load between a destination node and the source node |
| % Byte Load | Percentage of the byte load between a destination node and the source node |
| Pkts/s In | Packets per second received by a destination node from the source node |
| Pkts/s Out | Packets per second transmitted by a destination node to the source node |
| Bytes/s In | Bytes per second received by a destination node from the source node |
| Bytes/s Out | Bytes per second transmitted by a destination node to the source node |
| Pkts In | Number of packets received by a destination node from the source node since the view was opened |
| Pkts Out | Number of packets transmitted by a destination node to the source node since the view was opened |
| KBytes In | Total kilobytes received by a destination node from the source node since the view was opened |

| Statistic | Explanation |
|----------------|--|
| KBytes Out | Total kilobytes transmitted by a destination node to the source node since the view was opened |
| Protocols | Protocol packet types used by the destination node in this conversation |
| First Transmit | Date and time that the destination node first transmitted on the network since the traffic analysis agent was loaded |
| Last Transmit | Date and time that the destination node last transmitted since the traffic analysis agent was loaded |
| MAC Address | Physical (MAC) address of the destination node |

Monitoring Nodes for Inactivity

For segments on which at least one Traffic Analysis Agent for NetWare version 1.0 or later is installed, you can specify the nodes on the segment you want to monitor so that you are alerted if they become inactive. You can do this using the Monitor Nodes for Inactivity view.

Monitoring nodes for inactivity has the following advantages:

- ◆ You can monitor any node on the segment, regardless of the protocol the node uses.
- ◆ This feature does not impact network traffic because the traffic analysis agent does not poll the nodes to obtain their status.

To view a list of nodes monitored for inactivity:

- 1 In ConsoleOne, select a segment.
- 2 Click *View > Monitor Nodes for Inactivity*.

Another way to monitor connectivity is to specify the target in the Ping window and test the status of the specified node. The Connectivity Test window displays statistics that enable you to determine the status of the specified target. For details, see [Chapter 28, “Monitoring Services,” on page 1023](#).

By default, the poll interval for refreshing the Monitor Nodes for Inactivity view is zero seconds. You can configure the poll interval based on how often you want the view to be refreshed. For details, see [“Specifying the Poll Interval for Refreshing the Monitor Nodes for Inactivity View” on page 1088](#). You can also change the duration for the agent to verify the node before declaring it inactive. For details, see [“Specifying the Duration for the Agent to Determine if a Node Is Inactive” on page 1088](#).

IMPORTANT: You do not need to keep the Monitor Nodes for Inactivity view open or Novell ConsoleOne for the nodes to be monitored because the RMON agent is doing the monitoring, not Novell ConsoleOne. The Alarm Manager must be running to record an inactive node in the Alarm Report. If Novell ConsoleOne is not running, check for alarms after you restart it.

To monitor a node for inactivity:

- 1 In ConsoleOne, right-click a node or from any view that displays a list of nodes, then click *Monitor Nodes for Inactivity > Add*.

To disable a node from being monitored for inactivity:

- 1 In ConsoleOne, right-click the node that is monitored for inactivity then click *Monitor Nodes for Inactivity*, and then click *Delete*.

IMPORTANT: After the addition of any inactive node, if the NIC card of the node is changed, you will be able to see the node in the Monitor Node for Inactivity view but will not be able to delete it because of the change of MAC address.

Statistics displayed in the Monitor Nodes for Inactivity view are described in [Table 29-22](#):

Table 29-22 *Inactivity Statistics for the Monitor Nodes*

| Statistic | Explanation |
|-------------|--|
| Name | Displays a list of nodes that are being monitored for inactivity |
| MAC Address | Displays the MAC address of the network interface |
| Status | Displays the status of a node as active or inactive |

You can open the Monitor Nodes for Inactivity view to check the Status column any time Novell ConsoleOne is running. To do this, complete the following steps:

- 1 In ConsoleOne, select a segment.
- 2 Click *View > Monitor Nodes for Inactivity*.

The Status column displays if the selected node is active or inactive.

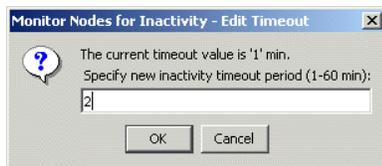
This Monitoring Nodes for Inactivity feature also allows you to add and monitor the nodes from the segments. To do this, follow these steps:

- 1 In ConsoleOne, right-click the segment, and select *Monitor Nodes For Inactivity view*.
- 2 Click  or select *File > Actions > Add Node(s)*.
- 3 In the Monitor Nodes for Inactivity - Add Node dialog box, select the nodes you want to add.

The nodes you add will be displayed in the view.

The Monitor Nodes for Inactivity module now sends alarms if the node you are monitoring is in the Off state or Timeout state. You can also change the timeout value from the Monitor Nodes for Inactivity view.

- 1 Click  or select *File > Actions > Edit Timeout*.



- 2 Specify the timeout period.
- 3 Click *OK*.

29.4.3 Capturing Packets

Novell ZENworks Server Management provides packet capture and decoding tools that help you analyze your network activity and identify the source of network problems. Capturing and decoding packets can help you troubleshoot network problems by giving you detailed information about what is actually happening on a segment.

Novell ConsoleOne can request packet capture on any monitored segment. Each RMON agent captures packets on the segment it monitors and stores information in its local buffer.

The following sections contain detailed information about capturing packets:

- ♦ “Defining a Capture Filter” on page 1063
- ♦ “Starting Packet Capture” on page 1065
- ♦ “Creating Simultaneous Packet Capture” on page 1065
- ♦ “Stopping Packet Capture” on page 1065
- ♦ “Restarting a Stopped Packet Capture” on page 1066
- ♦ “Saving and Viewing the Captured Packets” on page 1066

Defining a Capture Filter

Novell ZENworks Server Management provides a capture filter with default values you can use to capture packets using any Traffic Analysis agent. You can modify the values by defining a filter. For example, if you want to capture only NetWare packets sent by a certain node, you can define a filter to capture only those packets. As a result, the buffer has more space to store your selected packets. Once you define a filter setting it shall be saved for future uses.

When you specify a capture filter, you are specifying the packets to capture (include) in the buffer on the RMON agent, not the packets to exclude. When you specify both a node and a protocol, packets must meet both criteria to be captured. If you select more than one protocol family, packets can meet either protocol criteria to be captured.

To define a capture filter:

- 1 Click *File > Actions > Capture Packets*.
- 2 In the *Capture Name* text box, enter a name if you do not want to use the default name.
The capture name helps you keep track of multiple captures on the same segment.
- 3 In the *Capture Using* drop-down list, select the Traffic Analysis Agent that will be used to capture packets. Click *Find In Atlas* to select any Traffic Analysis Agent on the site. The Traffic Analysis Agent you have selected will be saved and used the next time.
- 4 Enter or select the source and destination nodes from the *Stations* box. You can also click the *Find Node* icon to select the node from the Find dialog box, an atlas component.
The *Stations* box displays a list of nodes on the segment from which the user can capture packets. You can select from Hardware, IP, or IPX stations.
If you choose ANY in both the source and destination node list, all packets sent by or received from any node are captured.
- 5 Select the direction of traffic flow between the nodes.
Click an arrow option from the drop-down list to specify the direction of the traffic flow. The available node and traffic flow directions are shown in the following table.

| Node | Arrow | Node | Effect |
|-------|-------|-------|---|
| node1 | <==> | node2 | Capture packets that node1 sends to node2 and packets that node2 sends to node1. |
| node1 | <==> | ANY | Capture packets that node1 sends to any node and packets that node1 receives from any node. This is equivalent to ANY <==> node1. |
| ANY | <==> | ANY | Capture all packets sent by or received from any node. |
| node1 | ==> | node2 | Capture packets that node1 sends to node2. This is equivalent to node2 <== node1. |
| node1 | ==> | ANY | Capture packets that node1 sends to any other node. This is equivalent to ANY <== node1. |
| node1 | <== | node2 | Capture packets that node2 sends to node1. This is equivalent to node2 ==> node1. |
| node1 | <== | ANY | Capture packets that any node sends to node1. This is equivalent to ANY<== node1. |

- 6** If you want to filter on protocols used, add the protocol suites you want to the *Selected* list box. To add a protocol to the *Selected* list box, select it from the *Available* list box > click *Add*.
or

To delete a protocol from the *Selected* list box, select it then click *Remove*.

All protocols are selected by default when you first use Novell ZENworks Server Management. If no protocols are listed in the *Selected* list box, all protocols are captured.

See [“Protocol Decodes Suite Supported by Novell ZENworks Server Management” on page 1041](#) for details about the protocol decoding support that Novell ZENworks Server Management provides.

- 7** Specify what kind of packets to capture on Ethernet, FDDI, or token ring segments.

The default statistics for the segments are listed in the following table.

| Segment Type | Available Statistics | Default Statistics |
|--------------|---|--------------------------------|
| Ethernet | Only good packets, only error packets, or both good and error packets. | Good packets and error packets |
| FDDI ring | All packets, LLC packets, MAC packets, or SMT packets. | All packets |
| Token ring | All packets, non-MAC packets, or MAC packets. MAC packets are used to manage the operation of the token ring. | All packets |

- 8** Specify whether to stop packet capture or to overwrite the oldest packets in the buffer with newer ones when the buffer is full.

Continuing packet capture means that a stop criteria does not exist and new packets will overwrite those already captured. You will need to manually stop packet capture if you select to overwrite the oldest packets.

- 9** Specify a buffer size.

Select a buffer size from the drop-down list or specify the size you want. The default buffer size is 128 KB.

The RMON agent will attempt to provide the buffer size requested. If not enough space is available in server memory for a large buffer, the RMON agent cannot create the requested size.

10 Select a slice size.

A slice specifies the maximum number of bytes of each packet, counting from the packet header, to keep in the buffer. This helps maximize the number of packets you can store in your buffer space, as well as reduce the load on the RMON agent to process captured packets. If you want to decode protocol header information, you need only 100 to 150 bytes. The rest is typically data that you need only if you suspect a data corruption problem. However, on certain very large packets, slicing can cause incorrect decodes by truncating information.

Your capture filter is now set up. If you decide not to capture packets, click the Cancel button.

Starting Packet Capture

To start packet capture:

- 1 Define a capture filter. See [“Defining a Capture Filter” on page 1063](#) for the procedure.
- 2 Click *OK* to apply the filter settings on the preferred RMON agent of the segment.
- 3 Click *Start* in the Capture Status dialog box.

When you start packet capture, the Start button in the Capture Status dialog box toggles to read Stop and the activity indicator reflects the capture buffer storage as it progresses. As packets that meet the filter criteria are captured, the capture buffer will begin to store the packet data, and a box below it will display the number of packets captured. The needle stops turning when the capture buffer is full.

Creating Simultaneous Packet Capture

You can create simultaneous packet captures by repeating the procedure you followed to start the first capture. This lets you set up and run captures with different capture criteria.

You can run a maximum of 20 packet captures with different capture criteria.

Stopping Packet Capture

When you set up a capture filter, you choose whether to stop packet capture when the capture buffer is full or to continue to capture packets but overwrite the oldest packets in the buffer.

By default, the packet capture will stop when the capture buffer is full. If you select to overwrite when the buffer is full, you must stop packet capture manually.

To stop packet capture manually, click the *Close* button in the Capture Status dialog box.

IMPORTANT: If you restart packet capture from the Packet Capture Setup window, the existing buffer is deleted and refreshed.

Restarting a Stopped Packet Capture

When the Packet Capture Setup window is open, you can start and stop capturing packets using the *Start/Stop* toggle button in the Capture Status dialog box. If Novell ZENworks Server Management is capturing packets, the button is labeled *Stop*; if it is not capturing packets, the button is labeled *Restart*. The RMON agent buffer is cleared when you restart.

Saving and Viewing the Captured Packets

You can save captured packets to a file and view as many files as you want, either while you are viewing a capture buffer or independently.

To view the saved packet capture files:

- 1 Click *Tools > View Packet File*.

The File Open dialog box is displayed.

- 2 Browse and select the packet capture file.

The `.tr1` file extension will be appended automatically.

29.4.4 Displaying Captured Packets

You can display and view decoded packets stored in the capture buffer from the Trace Display window by clicking the *View* button in the Capture Status dialog box. If you display this window while packets are being captured, capture automatically stops.

Novell ZENworks Server Management retrieves packet data from the RMON agent only as necessary for Novell ConsoleOne to decode and display the packets as you view them. This minimizes the amount of packet data transferred between the RMON agent and Novell ZENworks Server Management. If you prefer not to display all the packets you captured, you can create a display filter to display only a defined group of captured packets. For details, see [“Defining the Display Filter” on page 1069](#).

The following sections provide information on how you can view captured packets and perform trace display operations:

- ◆ [“Viewing Captured Packets” on page 1067](#)
- ◆ [“Filtering Packets for Display” on page 1068](#)
- ◆ [“Defining the Display Filter” on page 1069](#)
- ◆ [“Selecting and Decoding a Different Packet” on page 1071](#)
- ◆ [“Highlighting Protocol Fields and Hexadecimal Bytes” on page 1071](#)
- ◆ [“Saving Packet Files” on page 1071](#)
- ◆ [“Opening Packet Files” on page 1072](#)
- ◆ [“Printing Packets” on page 1072](#)

Novell ZENworks Server Management provides default settings based on which captured packets are displayed in the Trace Display window. To change the default values provided for displaying captured packet, see [“Choosing Options to Display a Captured Packet” on page 1085](#).

Viewing Captured Packets

You can use the Trace Display view to view the decoded packet capture information, the packet data in hexadecimal format, and a summary of the captured packets:

To view a captured packet:

- 1 In Novell ConsoleOne, select a node or a segment.
- 2 Click *File > Actions > Capture Packet*.
- 3 Capture packets using the capture filter of your choice. See [“Defining a Capture Filter” on page 1063](#) for details.
- 4 Click the *View* button in the Capture Status dialog box.

The Trace Display window contains three panes that display captured and decoded packets, as described in the following sections:

- ◆ [“Viewing the Packet Decode” on page 1067](#)
- ◆ [“Viewing Packet Data in Hexadecimal Format” on page 1067](#)
- ◆ [“Viewing a Summary of Captured Packets” on page 1068](#)

When you view packets initially, the first packet in the Summary pane is highlighted and selected. The contents of that packet are displayed in the Decode pane. If you select a different packet in the Summary pane, it is highlighted and the Decode pane displays its decoded contents.

You can change the size of the Trace Display panes by dragging the divider between windows.

Viewing the Packet Decode

The Decode pane displays detailed information about the contents of a selected packet. The packet contents are interpreted (decoded) and displayed by protocol fields.

By default, the Decode pane displays fully decoded packet data. You can configure the Trace Display window to display the decoded packets either as full protocol decodes or by one line per protocol layer. See [“Choosing Options to Display a Captured Packet” on page 1085](#) for details about how to change the default settings.

Viewing Packet Data in Hexadecimal Format

The Hexadecimal pane shows uninterpreted packet data in hexadecimal format. The ASCII or EBCDIC portion of the Hexadecimal pane (to the right) displays a dot for every hexadecimal byte that has no ASCII or EBCDIC equivalent.

The first column in the pane indicates the offset in hexadecimal bytes. The offset is the number of bytes counting from the beginning of the header. For example, the first three lines have the following offset:

- ◆ Hexadecimal 0—indicates zero offset
- ◆ Hexadecimal 10—indicates decimal 16 offset (16 bytes precede this)
- ◆ Hexadecimal 20—indicates decimal 32 offset (32 bytes precede this)

Regardless of whether you choose to display one-line decoded or fully decoded packets in the Decode pane, entire packets are displayed in the Hexadecimal pane. The Hexadecimal pane and the highlighting tool are especially helpful with the full-decode display when you are trying to associate

protocol fields with specific bytes in a packet. For details, see “[Highlighting Protocol Fields and Hexadecimal Bytes](#)” on page 1071.

Viewing a Summary of Captured Packets

The Summary pane gives you an overview of the conversation between the source and the destination nodes. You can select a packet in this pane for further decoding and display in the other panes. You can scroll the pane horizontally, and you can change the size and position of the columns in the pane.

Statistical information about the captured packets displayed by the Summary pane is described in [Table 29-23](#):

Table 29-23 *Statistical information of the captured packets as displayed by the Summary pane*

| Statistic | Explanation |
|------------------|--|
| No. | Numbers the packets in order of arrival at the traffic analysis agent. |
| Source | IP address, IPX address, or the physical (MAC) address of the node that sent the packet. Names are stored in the database. If no name is found in the database, the MAC address is displayed. |
| Destination | Node to which the packet was sent. The node is displayed as the IP address, IPX address, or the physical (MAC) address of the node. |
| Layer | Abbreviation of the highest protocol layer in the packet. It might display NCP for NetWare Core Protocol™ (NCP™) software, ether for the Ethernet data link layer, RTMP for the AppleTalk Routing Table Maintenance Protocol layer, or 802.2 for the IEEE 802.2 Logical Link Control layer. If you choose the full decode option, the Decode pane displays the full name of the protocol layer and all its fields. The Hexadecimal pane shows the entire packet. |
| Summary | Brief description of the contents of the highest protocol layer. |
| Error | Type of errors, if any, in the packet. This column is displayed only for Ethernet media. |
| Size | Number of bytes in the packet. Packet size always excludes the packet preamble and the CRC. |
| Absolute Time | Clock time on your computer when the packet arrived. |
| Interpacket Time | Time elapsed from the end of the preceding packet to the end of the current packet. |
| Relative Time | Time that elapsed since the arrival of the first packet still in the buffer. |

Filtering Packets for Display

After you have captured packets, you can apply a display filter to the capture buffer and view only the packets that interest you. You can filter on node names or addresses, protocol families or protocol layers, or contents of a selected field. This is useful in situations when, after you have captured packets, you realize there is a problem with a specific workstation and you want to display only the packets it has sent or received.

Display filtering requires the transfer of a portion of every captured packet from the RMON agent to Novell ConsoleOne. For large captures, this consumes time and network bandwidth. We recommend that you define very specific capture filters rather than filtering during display. However, subsequent filtering of the same capture does not result in additional data transfer from the traffic analysis agent because the data is already transferred to Novell ConsoleOne. Therefore, it is much quicker to filter the same packet capture a second time.

Display filters affect only the display; they do not change the capture buffer. All captured packets remain in the capture buffer and are available for viewing with a different display filter or without any display filter.

You can define a display filter in either of two ways:

- ♦ From the Trace Display window, click *View > Filter*.

The Display Filter dialog box is displayed. For details, see [“Defining the Display Filter” on page 1069](#).

- ♦ Double-click a packet in the Summary pane or double-click a selected protocol layer or field in the Decode or Hexadecimal pane.

A filter is set based on what you selected. You can also modify the filter information as needed. For details, see [“Point-and-Click Filtering” on page 1070](#).

Defining the Display Filter

Capture packets using the capture filter of your choice. See [“Defining a Capture Filter” on page 1063](#) for details. To define a display filter:

- 1** In Novell ConsoleOne, select a segment.
- 2** Click *File > Actions > Packet Capture*.
- 3** Click the *View* button in the Capture Status dialog box.
- 4** With the Trace Display window displayed and active, click *View > Filter*.
- 5** Select the nodes from the drop-down lists. You can select from IP, IPX or MAC address.
Alternatively, you can enter a node name or address in place of ANY in either or both of the drop-down list boxes.
- 6** Select the direction of the traffic flow from the arrow options available in the drop-down list.
- 7** To display all the packets of a specific protocol layer:
 - 7a** Double-click a protocol suite name from the list of protocols to display a list of all the protocols in the suite.
 - 7b** Scroll through the list to find the protocol you want.
 - 7c** Select the protocol.
- 8** To display all the packets that have the same contents in a specific field:
 - 8a** Enter the offset in hexadecimal bytes.
You can count the offset in the *Hexadecimal* pane when the packet is decoded, using the offset column for guidance. See [“Viewing Packet Data in Hexadecimal Format” on page 1067](#) for details.
 - 8b** Specify whether the offset is counted from the beginning of the packet or from the beginning of a protocol layer.

If you choose the protocol layer option, you must select a specific protocol in the *Protocol* box.

8c Enter the data that you want to include in the filter.

8d Specify the format in which you want the data to be displayed. Select from hexadecimal, ASCII, or EBCDIC format options.

You can also fill in the values using point-and-click filtering. See “**Point-and-Click Filtering**” on page 1070.

9 Click *OK*.

The dialog box closes and Novell ZENworks Server Management begins to select the required packets from the capture buffer.

If you have a large capture buffer, Novell ZENworks Server Management displays the initial packets that pass the filter. Novell ZENworks Server Management continues to filter in the background while you examine these packets.

The Summary pane shows the list of filtered packets that met the criteria in the display filter. You can view and decode them as described earlier in this section.

Point-and-Click Filtering

You can define a display filter using the point-and-click method by double-clicking a field in the Trace Display window.

To define a display filter using the point-and-click method:

1 To display only packets in one conversation (for example, between a node and a server), double-click a packet in that conversation in the Summary pane.

The Display Filter dialog box displays the source and destination of the selected packet. You can also modify the addresses, if needed. For example, you can change the destination address to ANY, the broadcast address, or a specific node address.

or

To display all the packets containing a specific protocol layer, double-click the protocol line in the *Decode* pane.

The Display Filter dialog box displays the protocol you selected.

or

To display all packets with the same contents as a specific field, double-click the field in the *Decode* pane.

The Display Filter dialog box displays the field, data, and type of data for the selected field.

or

To display all packets with the same content as a specific offset, click the field in the *Hexadecimal* pane.

The Display Filter dialog box displays the offset and the type of data for the selected field.

2 Click *OK*.

The dialog box closes and Novell ZENworks Server Management begins to select the packets from the capture buffer.

The Summary pane displays the list of packets that met the display filter criteria.

Selecting and Decoding a Different Packet

To select a different packet for decoding:

- 1 Select *View > Go To*.

You can also use the arrow keys on your keyboard to highlight a different packet.

- 2 Enter the packet number.

If the packet number specified is more than the total number of captured packets, an error message displays. If a display filter is set and the specified packet number has not passed the filter, then a packet closest to the specified packet is displayed.

Packets are retrieved from the RMON agent as you select their headers in the Summary pane using the mouse or the arrow keys. Using the Go To dialog box avoids transferring unwanted packet data from the RMON agent. Similarly, scrolling the Summary pane with the scroll button retrieves only the packet header data when creating the decode summary, whereas using the arrow keys retrieves all packet data.

Highlighting Protocol Fields and Hexadecimal Bytes

Novell ZENworks Server Management provides a highlighting tool that helps you associate protocol fields and hexadecimal bytes. Highlighting can be a useful training tool for new network managers who want to learn about protocol decoding.

You can use this tool in the following ways:

- ◆ Highlight a protocol layer in the Decode pane.
All bytes are highlighted in the selected protocol layer of the Hexadecimal pane.
- ◆ Click a field in any of the protocol layers in the Decode pane.
Associated bytes are highlighted in the Hexadecimal pane.
- ◆ Click hexadecimal bytes in the Hexadecimal pane.
All hexadecimal and ASCII or EBCDIC bytes of this field in the Hexadecimal pane are highlighted, and the associated field is highlighted in the Decode pane.
- ◆ Click ASCII or EBCDIC text in the Hexadecimal pane.
All hexadecimal and ASCII or EBCDIC bytes that belong to the field are highlighted in the Hexadecimal pane, and the associated field is highlighted in the Decode pane.

Saving Packet Files

You can save captured packets to a file and open the file later to analyze or print. When you save packets to a file, Novell ZENworks Server Management creates a binary file with the name you specify. You might want to save packets to a file in the following situations:

- ◆ To transfer the packets to another system or to send them for analysis.
- ◆ To apply a display filter to decoded captured packets so you can view only the packets that interest you. After you apply the display filter, you can save the filtered packets to a file.
- ◆ To compare packets saved from your buffer with other packets. You can either save the other packets, or view them from the capture buffer. You can view only one active capture buffer at a

time. However, after you have saved packets to a file, you can open as many files as you want, and simultaneously view a capture buffer, if desired.

Packet files are compatible with the Traffic Analysis Agent for Windows and earlier versions of ManageWise[®]. Hence, packets captured and saved using Traffic Analysis Agent for Windows can be viewed using Novell ZENworks Server Management.

To save captured packets to a file while viewing the capture buffer:

- 1 Click *File > Save As*.

The Save Filtered Packets or Save Unfiltered Packets dialog box is displayed, depending on whether you filtered your packets.

- 2 Enter the name in the *Filename* option.

The `.tr1` file extension is appended automatically.

- 3 Click *OK*.

IMPORTANT: Filter out the captured packets you want to save. (See “[Filtering Packets for Display](#)” on page 1068.) When you save packets, you save only those that pass the display filter. If you did not filter the display, all packets are saved.

Opening Packet Files

To open a packet file:

- 1 From the main menu of Novell ConsoleOne, click *Tools > View Packet File*.
- 2 Double-click the file you want to open.

Printing Packets

To print packets:

- 1 Open a Trace Display window, either by capturing packets or by opening a packet file.
- 2 Click *File > Print*.
- 3 Select the print options you want.

You can select the destination, format, and the packets you want to print.

- ♦ Choose whether to print to your default printer or to a file. If you choose a file, enter its name and specify whether the current packet data should overwrite the file or be appended to it.
- ♦ Choose whether you want a summary of the packet information, only the hexadecimal information, a full decode, or a brief decode. These formats correspond to the three panes described in “[Viewing Captured Packets](#)” on page 1067.
- ♦ Choose whether to print all packets, a range of packets, or only the filtered packets.

- 4 Click *OK*.

29.4.5 Analyzing Traffic Generated by Protocols in Your Network

Novell ZENworks Server Management lets you determine the distribution of protocols in your network and provides statistical information of the protocols discovered by the RMON2 agent in the network, as well as transport and application layers. You can also add supported and custom protocols to your network. Supported protocols are those that the RMON2 agent is able to decode and count the number of packets transmitted in your network using the protocol. Custom protocols are not supported by the RMON2 agent but are used by nodes in your network.

The following sections explain how you can use Novell ZENworks Server Management to manage protocols in your network:

- ♦ “[Displaying a List of Protocols Used in Your Network](#)” on page 1073
- ♦ “[Determining the Distribution of Protocols in a Segment](#)” on page 1075

Displaying a List of Protocols Used in Your Network

You can use the Protocol Directory property page to view a hierarchical representation of supported and custom protocols used in the network, transport, and application layers in your network. By default, the page displays the Protocol Directory Tree that displays a collapsed list of protocols. The protocols used in the data link layer are displayed at the top level. You can expand each protocol to display the list of supported and custom protocols under the selected protocol.

You can also use the Protocol Directory property page to add or delete the protocols supported by the RMON2 agent. For details, see “[Adding Supported Protocols to the Protocol Directory Tree](#)” on page 1074. The custom protocols that are used by the nodes in your network but are not supported by the RMON2 agent can also be added using the limited extensibility feature of RMON2. For details, see “[Adding Custom Protocols to a Supported Protocol Tree](#)” on page 1074. For details about the limited extensibility feature, see [RFC 2021 \(http://www.isi.edu/in-notes/rfc2021.txt\)](http://www.isi.edu/in-notes/rfc2021.txt).

For a selected protocol, you can specify the RMON2 groups you want the RMON2 agent to support. This will let you obtain the RMON2 details of the groups that you specify the agent to support. While adding the protocol, you can enable the agent support for the Host group, Matrix group, and Address Map group. The *Groups Supported* box in the lower portion of the property page indicates whether the agent support for the Host and Matrix groups in the network layer and application layer, and support for the Address Map group are enabled, disabled, or not supported for the selected protocol. You can configure the values displayed in the *Groups Supported* box.

The *Add* and *Remove* buttons are enabled only when you select a protocol in the Protocol Directory tree.

IMPORTANT: The Traffic Analysis Agent for NetWare and Traffic Analysis Agent for Windows do not support enabling of the Address Map, Host, and Matrix groups for protocols in the Protocol Directory.

To open the Protocol Directory property page:

- 1 In Novell ConsoleOne, click *RMON2* under Service within a node.
- 2 Click *File > Properties > Protocol Directory* tab.

Refer to the following sections:

- ◆ “Adding Supported Protocols to the Protocol Directory Tree” on page 1074
- ◆ “Adding Custom Protocols to a Supported Protocol Tree” on page 1074

Adding Supported Protocols to the Protocol Directory Tree

Supported protocols are those that the RMON2 agent is able to decode and count the number of packets transmitted in your network using the protocol.

Default values are provided for the parameters of protocols supported by the RMON2 agent. When you enter the name of a protocol, the default values are displayed if the protocol is supported.

To add a protocol to the Protocol Directory tree:

- 1 Open the Protocol Directory property page.
- 2 Select a protocol from the Protocol Directory tree.
- 3 Click *Add*.

The following table describes the parameters for a selected protocol.

IMPORTANT: The Protocol Name parameter cannot be configured. If you configure the port number or protocol code of a selected protocol, all child protocols of the selected protocol will be deleted.

| Parameter | Description |
|------------------|---|
| Protocol Name | Displays the name of the protocol. |
| Protocol ID | Displays the identifier for the protocol. Displays the port number for an application layer protocol or the protocol code for protocols in other layers. The protocol identifier is always a decimal value. |
| Description | Displays a short description of the selected protocol. |
| Groups Supported | Displays whether the agent support of the Address Map group, Host group, or Matrix group is enabled for the selected protocol. |

If the protocol name you enter or select from the *Protocol Name* list is supported by the RMON2 agent, the default parameters for the protocol are displayed in the appropriate fields of the Add Protocol dialog box. You cannot edit the parameters after you have added, if you do not want to use the default values.

- 4 Click *OK*.

The new protocol is added as a child protocol of the selected protocol. You cannot edit the parameters of the protocol you have added. You would need to delete the protocol and add the protocol again with different parameters.

Adding Custom Protocols to a Supported Protocol Tree

Custom protocols are those that are not supported by the RMON2 agent but are used by nodes in your network. If the RMON2 agent supports the limited extensibility feature of RMON2 for a selected protocol, you can add custom protocols under the selected protocol. See [RFC 2021 \(http://www.isi.edu/in-notes/rfc2021.txt\)](http://www.isi.edu/in-notes/rfc2021.txt) for more information. If the RMON2 agent does not support the

limited extensibility feature for a protocol, you cannot add custom protocols under that protocol. A custom protocol cannot have child protocols.

Because default values are not provided for custom protocols, you must enter the appropriate values if you are adding a protocol that is not supported by the RMON2 agent.

To add a custom protocol to the Protocol Directory tree:

- 1 Select a supported protocol from the Protocol Directory tree.
- 2 Click *Add*.
- 3 In the *Protocol Name* field, enter the name of the protocol.
- 4 In the *Protocol ID* field, enter the port number for an application layer protocol or a protocol code for protocols in other layers.

IMPORTANT: The port number or protocol code should be a decimal value.

- 5 From the *Groups Supported* box, select the groups you want the RMON2 agent to support for the protocol.

The custom protocol is added as a child protocol of the supported protocol.

To remove a protocol from the Protocol Directory tree:

- 1 Select a protocol from the Protocol Directory tree.
- 2 Click *Remove*.

IMPORTANT: If you remove a protocol that has child protocols, all the child protocols are also removed from the Protocol Directory tree.

Determining the Distribution of Protocols in a Segment

Novell ZENworks Server Management lets you determine the distribution of protocols discovered by the RMON2 agent. You can use the information displayed in this view to analyze the traffic in your network and to troubleshoot network problems. Use the Protocol Directory property page to add, delete, or edit a protocol. See [“Adding Supported Protocols to the Protocol Directory Tree” on page 1074](#) and [“Adding Custom Protocols to a Supported Protocol Tree” on page 1074](#) for details.

The distribution of protocols discovered by the RMON2 agent is displayed in the Protocol Distribution view, based on the layer in which the protocols are discovered.

To view the distribution of protocols in the selected segment:

- 1 In Novell ConsoleOne, select a segment.
- 2 Click *View > Protocol Distribution*.

The view displays the following three tables that list the protocols discovered in the network:

- ♦ Network layer table
- ♦ Transport layer table
- ♦ Application layer table

The protocols discovered by the RMON2 agent are placed in the appropriate table in the Protocol Distribution view depending on the layer in which they were discovered. Each table displays protocol statistics that are updated every 15 seconds.

Table 29-24 describes the protocol statistics displayed in the Protocol Distribution view:

Table 29-24 Protocol statistics displayed in the Protocol Distribution view

| Statistic | Description |
|---------------|---|
| Protocol Name | The name of the protocol |
| Packets/s | The average number of packets transmitted per second using the protocol discovered by the agent on the monitored segment |
| Bytes/s | The average number of bytes transmitted per second using the protocol discovered by the agent discovered on the monitored segment |
| Packet Rate % | The percentage of packets transmitted using the protocol; this is relative to the total percentage of packets transmitted using all protocols discovered by the agent |
| Byte Rate % | The percentage of bytes transmitted using the protocol; this is relative to the total percentage of bytes transmitted using all protocols discovered by the agent |

IMPORTANT: Only one entry of each protocol is displayed in the Protocol Distribution view. Consolidated statistics are displayed for a supported protocol in more than one protocol suite.

29.4.6 Analyzing Traffic on Switches

Novell ZENworks Server Management provides statistical information about ports in a monitored switch and a list of nodes connected to each port in your switched network. This information is displayed in the Unified Port Traffic view. You can use the view to determine the load on the desktop and workgroup switches in your switched network. When only one node can be connected to each port in a switch, the switch is known as a desktop switch. When one port of a switch is connected to a connecting device to which more than one node is connected, the switch is called a Workgroup switch.

Ports and nodes connected to ports of a switch can be monitored using an embedded RMON agent or external RMON agent. An embedded RMON agent is installed on the port of a switch. An external RMON agent is installed on a node connected to a switch.

The following sections explain how you can obtain information about switch ports and nodes connected to ports in your switched network:

- ◆ [“Viewing Statistics for Ports in a Switch” on page 1076](#)
- ◆ [“Viewing the Summarized Switch Information” on page 1077](#)

Viewing Statistics for Ports in a Switch

You can use the Unified Port Traffic view to obtain statistical information about every switch port in your network. The view also displays a drop-down list of nodes connected to each port. The information displayed in this view is useful if you want to troubleshoot a port.

The Unified Port Traffic view displays a list of nodes connected to ports on the switch and statistics for each port. You can view Ethernet specific statistics for Ethernet ports on a switch. Statistics specific to FDDI and token ring ports are not displayed with this version of Novell ZENworks Server Management, although general port statistics are displayed for all ports on a switch

regardless of the media type. You can choose to display all statistics or configure the Unified Port Traffic view to display selected statistics. For details, see “[Choosing Statistics to Display in the Unified Port Traffic View](#)” on page 1084.

To display the statistics of ports in a switch:

- 1 In Novell ConsoleOne, select *Switch/Bridge* under *Services* within a switch.
- 2 Click *View > Port Traffic*.

Viewing the Summarized Switch Information

The Switch Summary view provides brief information about a selected switch. You can view static information about a selected switch and information about alarms generated on the switch. You can also determine the packets and broadcasts received by the switch per second.

To view the summarized switch information:

- 1 In Novell ConsoleOne, select *Switch/Bridge* under *Services* within a switch.
- 2 Click *View > Switch Summary*.

The Switch Summary view displays static information about a selected switch, as described in [Table 29-25](#):

Table 29-25 *Switch static information displayed in the Switch Summary view*

| Statistic | Explanation |
|---------------------------------|--|
| Vendor | Name of the switch vendor |
| Switch Type | Type of switch: Transparent or Source Route |
| Number of Ports Active | Number of active ports on the switch |
| Forwarding Table Overflow Count | Number of times the forwarding table has exceeded its capacity |
| Up Time | Time since the switch was last rebooted |
| Number of Ports Present | Number of ports present on the selected switch |
| Number of MAC Addresses Learned | Number of MAC addresses dynamically discovered by the switch |

The Switch Summary view displays information about alarms generated on a selected switch, as described in [Table 29-26](#):

Table 29-26 *Information about alarms generated on a selected switch as displayed in the Switch Summary view*

| Statistic | Explanation |
|-----------|--|
| Severity | Severity level attributed to the trap. |

| Statistic | Explanation |
|---------------|--|
| From | Network address of the device that sent the alarm to the alarm management system. |
| Owner | Segment or device affected by the alarm. |
| Summary | Summary of the event, often including the name or address of the object affected by the alarm. |
| Received Time | Date and time when the alarm management system received the alarm. |
| Type | Generic description of the alarm. For example, Volume out of disk space. |
| Category | Displays the category of the alarm based on the MIB that defines the trap-type objects. The category is directly related to the MIBs included in the management server MIB pool. For example, the category for NetWare servers is based on the NetWare Server Alarm MIB. |

The Switch Summary view displays dynamic information about a selected switch, as described in [Table 29-27](#):

Table 29-27 *Dynamic information about a selected switch as displayed in the Switch Summary view*

| Statistics | Explanation |
|------------------------|---|
| Switch Load (pkts/sec) | The load on the switch based on packets received by the switch per second |
| Frames Dropped/sec | The number of received packets discarded per minute |
| Broadcasts/sec | The number of broadcasts received by the switch from the nodes connected to ports of the switch |

29.5 Optimizing Traffic Analysis

The tools provided by Novell ZENworks Server Management to analyze your network performance have default settings. You can change the default settings of various views to display only the information you require.

The following sections provide information about how you can configure the Novell ZENworks Server Management tools to suit your networking environment:

- ◆ [Section 29.5.1, “Choosing Options to Display Stations on a Segment,” on page 1079](#)
- ◆ [Section 29.5.2, “Choosing Options to Display Trend Statistics,” on page 1080](#)
- ◆ [Section 29.5.3, “Choosing Options to Display the Top Nodes Graph,” on page 1083](#)
- ◆ [Section 29.5.4, “Choosing Statistics to Display in the Unified Port Traffic View,” on page 1084](#)
- ◆ [Section 29.5.5, “Choosing Options to Display a Captured Packet,” on page 1085](#)
- ◆ [Section 29.5.6, “Configuring Alarm Options from the Set Alarm Dialog Box,” on page 1085](#)
- ◆ [Section 29.5.7, “Configuring the Monitor Nodes for Inactivity View,” on page 1088](#)

29.5.1 Choosing Options to Display Stations on a Segment

You can configure the Stations view to display only the top 20 nodes or all nodes on the monitored segment. You can also choose the statistic based on which you want to display the top 20 nodes.

The following configuring options are available:

- ◆ “[Displaying Statistics for All Nodes on a Segment](#)” on page 1079
- ◆ “[Displaying Statistics for the Top 20 Nodes on a Segment](#)” on page 1079
- ◆ “[Choosing a Statistic Based on Which Top 20 Nodes Are Displayed](#)” on page 1079

Displaying Statistics for All Nodes on a Segment

To display statistics for all nodes on a segment:

- 1 In Novell ConsoleOne, select a segment.
- 2 Click *View > Stations*.
To display all nodes on a segment, more time is required and more network traffic is generated.
- 3 From the Stations view, click *View > Show All Stations*.

Displaying Statistics for the Top 20 Nodes on a Segment

To display statistics for the top 20 nodes on a segment:

- 1 In Novell ConsoleOne, select a segment.
- 2 Click *View > Stations*.
- 3 From the Stations view, click *View > Show Top N Stations*.

Choosing a Statistic Based on Which Top 20 Nodes Are Displayed

Packets out per second is the default statistic based on which top 20 nodes are displayed in the Stations view.

To choose a different statistic based on which you want the top 20 nodes to be displayed, do either of the following:

- 1 From the Stations view, click *View*, select *Show Top N Stations*, then choose a statistic from the list of statistics displayed.
- 2 Click the *Top Nodes Statistics* drop-down box in the toolbar of the *Stations* view, then choose a statistic from those displayed.

The available statistics are described in [Table 29-28](#):

Table 29-28 *Top Nodes Statistics*

| Statistic | Explanation |
|---------------|--|
| Packets/s In | Packets per second received by a node |
| Packets/s Out | Packets per second transmitted by a node |
| Bytes/s In | Bytes per second received by a node |

| Statistic | Explanation |
|--------------|---|
| Bytes/s Out | Bytes per second transmitted by a node |
| Errors/s | Errors per second transmitted by a node |
| Broadcasts/s | Broadcast packets per second transmitted by a node |
| Multicasts/s | Multicast packets per second transmitted by a node (packets transmitted to a specific group of nodes) |

If you close the Stations view after changing the default settings, you will be prompted to save the changes made to the default settings. If you want the Stations view to be displayed based on the statistic you chose, you can save the setting. The next time you open Novell ConsoleOne and launch the Stations view, you will be able to view the nodes on the monitored segment based on the statistic you specified.

29.5.2 Choosing Options to Display Trend Statistics

You can change the default settings based on which the segment performance trends are displayed in the Segment Trends view.

The following configuration options are available:

- ◆ [“Choosing Statistics Based on Which Trend is Displayed” on page 1080](#)
- ◆ [“Setting the Time-Scale Options” on page 1082](#)

Choosing Statistics Based on Which Trend is Displayed

To change the statistics based on which segment performance trend is displayed:

- 1 Click the *Profile* button in the Segment Trends view.
- 2 Select a profile from the *Select Profile* list.

The default profile will display a trend with statistical information of total packets, good packets, and error packets on the monitored segment.

If you choose not to use the profiles listed in the *Select Profile* list, you can select the required statistics from the *Select Statistics* list.

The statistics list lets you examine the Ethernet, FDDI, and token ring statistics described in [Table 29-29](#):

Table 29-29 Ethernet, FDDI, and token ring statistics

| Statistic | Media Support | Explanation |
|--------------------------|---------------|--|
| Abort Delimiter Errors/s | Token ring | Average number of abort delimiter errors observed per second. This error indicates that a node aborts a transmission. |
| AC Errors/s | Token ring | Average number of AC errors observed per second. This error is reported when an intended recipient of a packet fails to mark it as received or flags an error on it. |

| Statistic | Media Support | Explanation |
|----------------------------|----------------------------|--|
| Beacons | FDDI and token ring | Average number of beacons per second observed in the sampling interval. A station transmits these packets when it detects a hard failure upstream. |
| Broadcast Packets/s | Ethernet, FDDI, token ring | Number of broadcast packets per second. |
| Burst Errors/s | Token ring | Average number of burst errors observed per second. This error indicates that a node detects the absence of transitions for the required time. |
| Claim Tokens/s | FDDI ring | Average number of times that the ring enters the claim token state from the normal ring state or ring purge state per second. |
| CRC/Alignment Errors/s | Ethernet and FDDI ring | Number of cyclic redundancy check (CRC)/alignment errors per second. |
| Echo Pkts/s | FDDI ring | Average number of echo frames received on the network per second. |
| Elasticity Buffer Errors/s | FDDI ring | Average number of elasticity buffer overflow errors reported by this station per second. This is due to the difference in the clock frequency between the transmitting and receiving stations. |
| Error Packets/s | Ethernet | Number of error packets per second. |
| Fragments/s | Ethernet | Number of fragments per second. |
| Frame Copied Errors/s | FDDI ring | Average number of frame copied error frames reported per second by the station. |
| Frequency Errors/s | Token ring | Average number of frequency errors observed per second. This error indicates that a token ring clock on a node differs too much from the clock on the active monitor. |
| Good Packets/s | Ethernet | Number of good packets per second. |
| Internal Errors/s | Token ring | Average number of internal errors observed per second. These errors generally indicate a network board failure. |
| Jabbers/s | Ethernet | Number of jabbers per second. |
| Line Errors/s | Token ring | Average number of line errors observed per second. These packets are of valid size but have a faulty Frame Check Sequence (FCS) and do not end on an 8-bit boundary. |
| Lost Frames/s | FDDI and token ring | Average number of lost frame errors on the network observed per second. |
| Monitor Contentions/s | Token ring | Average number of monitor contentions observed per second; these packets are transmitted by all active nodes when no active monitor is detected on the ring. |

| Statistic | Media Support | Explanation |
|-----------------------------|--------------------------------|---|
| Multicast Packets/s | Ethernet, FDDI, and token ring | Number of multicast packets per second. |
| Oversize Packets/s | Ethernet | Number of oversize packets per second. |
| Packets | FDDI and token ring | Average number of packets observed per second in the sampling interval. |
| Receive Congestion Errors/s | Token ring | Average number of receive congestion errors observed per second. This error indicates that a node recognizes a frame addressed to its address, but has no available buffer space. |
| Ring Wraps/s | FDDI ring | Average number of times a wraparound condition has been detected at this interface per second. This entry does not indicate the number of times the ring has actually wrapped around. It only indicates the number of times the ring has wrapped around this physical path. |
| Token Errors/s | Token ring | Average number of token errors observed per second. This error indicates that a token is corrupted or the active monitor did not see a new frame in the required amount of time. |
| Total Bytes/s | Ethernet | Average number of total bytes per second. |
| Total Packets/s | Ethernet | Average number of total packets per second. |
| Undersize Packets/s | Ethernet | Number of undersize packets per second. |
| Unicast Packets/s | Ethernet | Number of unicast packets per second. |
| Utilization% | Ethernet, FDDI, and token ring | Percentage of maximum network capacity used by all packets in the sampling interval. |

If you close the Segment Trends view after changing the default statistics based on which trend is displayed, you will be prompted to save the changes made to the default settings. If you want the segment performance trend to be displayed based on the profile or statistics you chose, you can save the settings that you define. The next time you open Novell ConsoleOne and launch the Segment Trends view, you will be able to view the trend based on the profile or statistics you defined.

Setting the Time-Scale Options

The segment performance trend is updated once every minute. You can set a different time scale based on which you want to update a graph. Select from the following time-scale options:

- ◆ Real Time
- ◆ One Hour
- ◆ One Day
- ◆ One Week
- ◆ One Month
- ◆ One Year

TIP: If you close the Segment Trends view after changing the default time-scale option based on which trend is displayed, you will be prompted to save the changes made to the default settings. If you do not want the trend to be updated in real time, you can save the time-scale setting you choose. The next time you open Novell ConsoleOne and launch the Segment Trends view, the trend will be updated based on the time-scale option you selected.

29.5.3 Choosing Options to Display the Top Nodes Graph

You can configure the Segment Dashboard view to display or disable the top nodes graph. For details, see [“Viewing the Graph of the Top Nodes on a Monitored Segment” on page 1050](#). The top nodes graph is displayed in the lower portion of the Segment Dashboard view. Packets out per second is the default statistic based on which the graph is displayed. You can choose a different statistic based on which you want the graph to be displayed.

The following configuring options are available:

- ◆ [“Displaying the Top Nodes Graph in the Segment Dashboard View” on page 1083](#)
- ◆ [“Choosing the Statistic Based on Which Top Nodes Graph Is Displayed” on page 1083](#)
- ◆ [“Disabling the Top Nodes Graph in the Segment Dashboard View” on page 1084](#)

Displaying the Top Nodes Graph in the Segment Dashboard View

To display the top nodes graph in the Segment Dashboard view:

- 1 From the Segment Dashboard view, click *View > Show Top N Graph*.

Choosing the Statistic Based on Which Top Nodes Graph Is Displayed

To display the top nodes graph based on a different statistic, do either of the following from the Segment Dashboard view:

- 1 Click *View > Show Top N Graph*, then choose a statistic.
- 2 Click the *Top Nodes Statistics* drop-down box in the toolbar of the *Segment Dashboard* view > select a statistic.

The statistics are described in [Table 29-39](#):

Table 29-30 *Top Nodes Statistics*

| Statistic | Explanation |
|----------------|--|
| Broadcasts/min | Broadcast packets per minute transmitted by a node |
| Bytes/s in | Bytes per second received by a node |
| Bytes/s out | Bytes per second transmitted by a node |
| Errors/min | Errors per minute transmitted by a node |
| Packets/s in | Packets per second received by a node |
| Packets/s out | Packets per second transmitted by a node |

| Statistic | Explanation |
|----------------|--|
| Multicasts/min | Multicast packets per minute transmitted by a node |

IMPORTANT: Errors per minute, broadcasts per minute, and multicasts per minute are updated every 60 seconds rather than every 5 seconds.

Disabling the Top Nodes Graph in the Segment Dashboard View

To disable the top nodes graph in the Segment Dashboard view:

- 1 From the Segment Dashboard view, click *View > Disable Top N Graph*.

29.5.4 Choosing Statistics to Display in the Unified Port Traffic View

Novell ZENworks Server Management provides statistics for each port on the switch. You can view port statistics and a list of nodes connected to each port using the Unified Port Traffic view. You can view Ethernet-specific statistics for Ethernet ports on a switch. Although statistics specific to FDDI and token ring ports will not be displayed with this version of Novell ZENworks Server Management, general port statistics are displayed for all ports on a switch regardless of the media type. For details, see “[Viewing Statistics for Ports in a Switch](#)” on page 1076. You can choose to display only the selected statistics in the Unified Port Traffic view.

To select statistics to be displayed in the Unified Port Traffic view:

- 1 From the Unified Port Traffic view, click *View > Settings*.
- 2 Click the statistics from the *Available Columns* list > click *Add*.

Table 29-31 describes the general port statistics displayed for a port, regardless of the media type of the port:

Table 29-31 General port statistics displayed for a port

| Statistic | Explanation |
|------------------|---|
| Frames In/sec | Number of frames received by the port per second. |
| Frames Out/sec | Number of frames sent by port per second. |
| Port Link Status | Displays if the port is active or inactive. If the port is active, it can transmit and receive packets. |
| Speed | The speed at which packets are transmitted or received by the port. |
| Media Type | Media type of the selected port. |
| Local Traffic | Rate of traffic going towards nodes on the same port. |

Table 29-32 describes the Ethernet-specific statistics displayed for an Ethernet port in addition to the general port statistics listed above:

Table 29-32 Ethernet-specific statistics displayed for an Ethernet port

| Statistic | Explanation |
|-----------------|---|
| Collisions/sec | Number of collisions per second |
| Utilization | Percentage of maximum network capacity currently consumed by packet traffic on the port |
| Broadcasts/sec | Number of broadcast packets per second currently received and sent by the port |
| Multicasts/sec | Multicast packets per second received and sent by the port |
| Packets/sec | Number of packets per second received and sent by the port |
| CRC Align Error | Total number of line errors reported by the port |
| Oversize Pkts | Number of oversize packets received and sent by the port |

29.5.5 Choosing Options to Display a Captured Packet

Novell ZENworks Server Management provides default settings to display a captured packet in the Trace Display window.

To change the default settings and display the trace differently:

- 1 Open the Trace Display window.
- 2 From the Trace Display menu, click *View > Options*.
- 3 Select how you want to display the decoded packet.
 - ♦ Full Protocol Decode: Provides information about each field in each protocol layer in a selected packet. This is the default decoding.
 - ♦ One Line Per Protocol Layer: Provides a line of information for each protocol layer of a selected packet.
- 4 Select the level at which you want to display the initial highlight position.
 - ♦ At Highest Protocol Layer: Places the initial highlighting at the highest protocol layer in a packet. This is the default.
 - ♦ At Packet Header: Places the initial highlighting at the packet header.
- 5 Select the format in which you want to display the decoded packet.
 - ♦ ASCII: Displays the hex data in ASCII format. This is the default.
 - ♦ EBCDIC: Displays the hex data in EBCDIC format.

29.5.6 Configuring Alarm Options from the Set Alarm Dialog Box

Novell ZENworks Server Management provides default alarm threshold values for a segment. You can set threshold values for various error conditions on Ethernet, FDDI, and token ring segments to eliminate the need to constantly monitor the segments.

When a segment alarm is enabled, the RMON agent monitors the segment based on the alarm threshold settings. If the configured threshold value is exceeded, the RMON agent sends a trap to the management server, which forwards it to Novell ConsoleOne.

You should change the default values for alarm thresholds as appropriate for your organization. You can determine the appropriate value by observing average and peak traffic levels on your network using the Segment Trends view. For details, see [“Analyzing Trend Data for a Segment” on page 1050](#). You can do this as a part of creating a baseline of typical segment activity on your network.

To set an alarm threshold for a segment:

- 1 In Novell ConsoleOne, select a segment.
- 2 Click *File > Properties > the Segment Alarms* tab.
- 3 Select a segment statistic > click *Edit*.
- 4 Click *Enable* to enable the alarms set for the monitored segment.

When you click *Enable*, the text fields and the *Default* button will be enabled. However, if the default threshold values are not found, the *Default* button will not be enabled.

- 5 Enter the threshold value.
- 6 Specify the sampling time interval.

The RMON agent uses the sampling time interval to average the statistic to determine whether the alarm threshold was exceeded.

TIP: You can also use the Segment Dashboard view to define alarm threshold values for segment statistics. For details, see [“Defining Alarm Thresholds for Statistics Displayed in the Segment Dashboard View” on page 1050](#).

Table 29-33 describes the alarm statistics that Novell ZENworks Server Management tracks for Ethernet, FDDI, and token ring segments:

Table 29-33 Alarm statistics used by ZENworks Server Management to track for Ethernet, FDDI, and token ring segments

| Statistic | Media Support | Explanation |
|--------------|--------------------------------|---|
| Abort Errors | Token ring | Average number of abort errors observed per second in the sampling interval. These errors resemble line errors, but occur in the middle of a transmission. |
| AC Errors | Token ring | Average number of Address Recognition (and Frame Copied) errors observed per second in the sampling interval. This error is reported when an intended recipient of a packet fails to mark it as received or flags an error on it. |
| Beacons | FDDI and token ring | Average number of beacons per second observed in the sampling interval. A station transmits these packets when it detects a hard failure upstream. |
| Broadcasts | Ethernet, FDDI, and token ring | Average number of packets per second sent to the broadcast address FF-FF-FF-FF-FF-FF. Broadcast messages typically consist of general requests for information or transmission of status information to all stations. |
| Burst Errors | Token ring | Average number of burst errors observed per second in the sampling interval. A burst error is caused by a lack of signal transitions between stations for a short period of time. |

| Statistic | Media Support | Explanation |
|----------------------------|--------------------------------|---|
| Claim Tokens | FDDI ring | Average number of times that the ring enters the claim token state from the normal ring state or ring purge state per second. |
| Congestion Errors | Token ring | Average number of congestion errors observed per second in the sampling interval. The receiving station runs out of buffer space to store the packet. |
| CRC Errors | Ethernet and FDDI ring | Average number of CRC errors observed per second in the sampling interval. These packets are of valid size but have a faulty FCS. |
| Echo Pkts | FDDI ring | Average number of echo frames received on the network per second. |
| Elasticity Buffer Errors/s | FDDI ring | Average number of elasticity buffer overflow errors reported per second by this station. This is due to the difference in the clock frequency of the transmitting and receiving stations. |
| Fragments | Ethernet | Average number of fragments observed per second in the sampling interval. Fragments are packets that contain fewer than 64 bytes and have a faulty FCS. They are typically a result of collisions. |
| Frame Copied Errors | FDDI and Token ring | Average number of frame copied errors observed per second in the sampling interval. This error indicates that a station has detected that another station accepted a packet addressed to the first station. |
| Frequency Errors | Token ring | Average number of frequency errors observed per second in the sampling interval. This error indicates that a token ring clock on a station differs from the clock on the active monitor. |
| Internal Errors | Token ring | Average number of internal errors observed per second in the sampling interval. These errors generally indicate a network adapter board failure. |
| Jabbers | Ethernet | Average number of jabber packets observed per second in the sampling interval. A jabber consists of packets that contain more than 1518 bytes and have a faulty FCS. |
| Line Errors | Token ring | Average number of line errors observed per second in the sampling interval. These packets are of legal size but have a faulty FCS and do not end on an 8-bit boundary. |
| Lost Frames | FDDI and token ring | Total number of lost frame errors received on the network. A lost frame error indicates that the end delimiter of a frame is lost in the network. |
| Monitor Contentions | Token ring | Average number of monitor contentions observed per second in the sampling interval. These packets are transmitted when no active monitor is detected on the ring. |
| Multicasts | Ethernet, FDDI, and token ring | Average number of packets per second sent to multicast addresses. |
| Oversize | Ethernet | Average number of oversized packets observed per second in the sampling interval. Oversized packets contain more than 1518 bytes, including the FCS. |
| Packets | Ethernet, FDDI, and token ring | Total number of packets observed per second in the sampling interval. |

| Statistic | Media Support | Explanation |
|----------------|--------------------------------|--|
| Ring Wraps/s | FDDI ring | Average number of times a wraparound condition has been detected at this interface per second. This entry does not indicate the number of times that the ring has actually wrapped around. It only indicates the number of times the ring has wrapped around this physical path. |
| Token Errors | Token ring | Average number of token errors observed per second in the sampling interval. This error indicates that a token is corrupted or the active monitor did not detect a new frame transmitted during the current sampling interval. |
| Undersize | Ethernet | Average number of undersized packets observed per second in the sampling interval. Undersized errors are shorter than 64 bytes. |
| Utilization(%) | Ethernet, FDDI, and token ring | Percentage of maximum network capacity used by all packets in the sampling interval. |

When you have set the appropriate threshold values for the segments in your network, you can use the Save As Default button on the Segment Alarms property page to save the values you defined as the default values. However, the default threshold values provided by Novell ZENworks Server Management will not be available after you apply the new values.

29.5.7 Configuring the Monitor Nodes for Inactivity View

By default, the poll interval for refreshing the Monitor Nodes for Inactivity view is zero seconds. You can configure the poll interval based on which you want the view to be refreshed. The agent monitoring nodes on a monitored segment declares a node as inactive after verifying it for a specified period of time. You can change the time duration for the agent to verify the node before declaring it inactive.

The following configuring options are available:

- ◆ [“Specifying the Poll Interval for Refreshing the Monitor Nodes for Inactivity View” on page 1088](#)
- ◆ [“Specifying the Duration for the Agent to Determine if a Node Is Inactive” on page 1088](#)

Specifying the Poll Interval for Refreshing the Monitor Nodes for Inactivity View

You can modify the `PollInterval` parameter in the `lsmparameters.properties` file to specify the poll interval for refreshing the Monitor Nodes for Inactivity view.

To specify a poll interval for refreshing the Monitor Nodes for Inactivity view:

- 1 Open the `lsmparameters.properties` file located in the `operating_system_drive\install\novell consoleone\bin` directory.
- 2 Specify a value for the `PollInterval` parameter.
The `PollInterval` value should be a positive value, in seconds. The default value is zero (0) seconds.

Specifying the Duration for the Agent to Determine if a Node Is Inactive

When a selected node becomes inactive, the agent monitoring the node verifies the state of the node for one minute before declaring it inactive. You can modify the `HostTimeout` parameter in the

`lsmparameters.properties` file to change the duration for the agent to verify the selected node before declaring it inactive. The agent verifies the inactive node for the specified period of time before declaring it inactive.

To change the duration for the agent to verify a node before declaring it inactive:

- 1 Open the `lsmparameters.properties` file located in the `operating_system_drive\install\novell consoleone\bin` directory.

- 2 Specify a value for the `HostTimeOut` parameter.

The `HostTimeOut` value should be a positive value, in minutes. The default value is one (1) minute.

29.6 Understanding the Traffic Analysis Agents

Traffic Analysis agents enable you to monitor a heterogeneous LAN environment comprised of Ethernet, FDDI, and token ring segments from the easy-to-use Novell ZENworks Server Management interface.

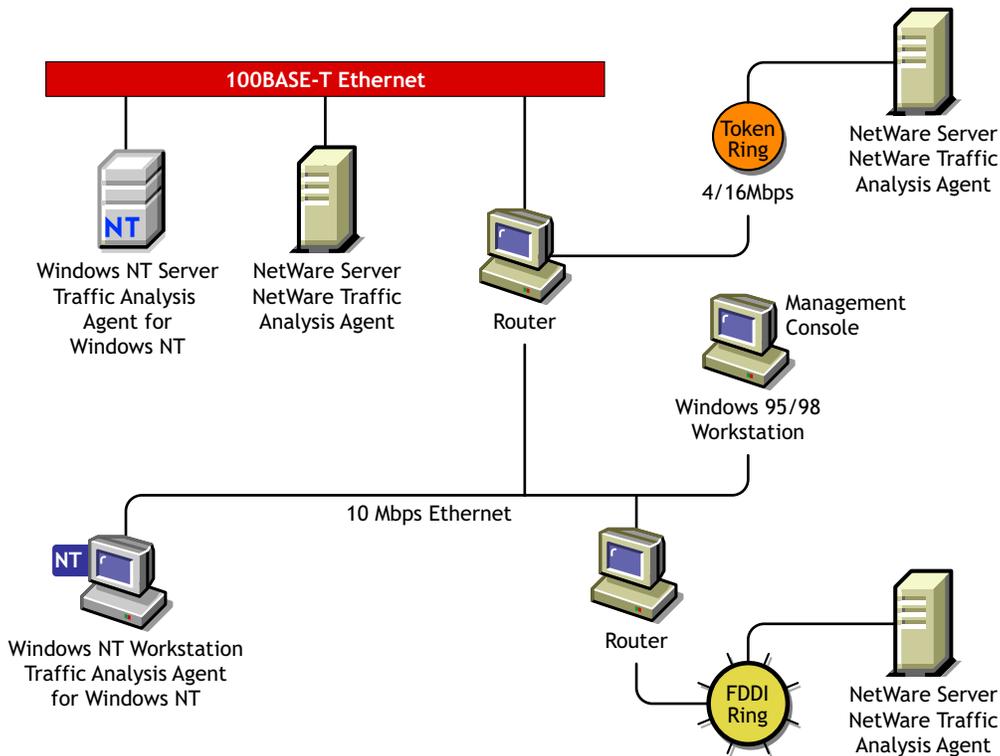
Traffic Analysis agents are RMON agents that can run on a NetWare server, Windows server, or a Windows workstation. They implement a set of functionality defined by the RMON MIB ([RFC 1757](http://www.isi.edu/in-notes/rfc1757.txt) (<http://www.isi.edu/in-notes/rfc1757.txt>)). These agents collect information about activity on your network and make it available to Novell ConsoleOne via SNMP.

The following functionality is provided by the Traffic Analysis Agents:

- ◆ Monitor the performance of segments and provide vital network statistical information to Novell ConsoleOne
- ◆ Make it easy to set alarm thresholds for proactive network management
- ◆ Capture all packets or selected packets to help you diagnose and resolve problems on the monitored networks
- ◆ Monitor multiple network segments including the Symmetric Multi-Processing (SMP) architecture
- ◆ Monitor network segments for problems, such as high network utilization and communication errors
- ◆ Track dynamic IP address assignments from the DHCP server to the nodes on the network
- ◆ Store data to display real-time trends (hourly) and historical trends (daily, weekly, monthly, and yearly) for statistics such as Total Bytes, Total Packets, Good Packets, Error Packets, and so forth
- ◆ Monitor nodes for inactivity, so that you are alerted if the monitored nodes becomes inactive

Figure 29-7 illustrates the functionality of traffic analysis agents:

Figure 29-7 Traffic Analysis Agents



Novell ZENworks Server Management includes the following traffic analysis agents:

- ♦ Traffic Analysis Agent for NetWare.
For details, see [Section 29.7, “Using the Traffic Analysis Agent for NetWare,” on page 1090](#).
- ♦ Traffic Analysis Agent for Windows.
For details, see [Section 29.8, “Using the Traffic Analysis Agent for Windows,” on page 1104](#).

The Novell ZENworks Server Management traffic analysis agents are RMON Plus agents. For details, see [“Functionality of RMON Plus Agents” on page 1035](#). These agents also implement the first two RMON2 groups. The first RMON2 group is the Protocol Directory group, which provides a table of protocols for which the agent will monitor and maintain statistics. The second RMON2 group is the Protocol Distribution group, which provides a table of statistics for each protocol in the directory. For details, see [“Functionality of RMON2 Agents” on page 1037](#).

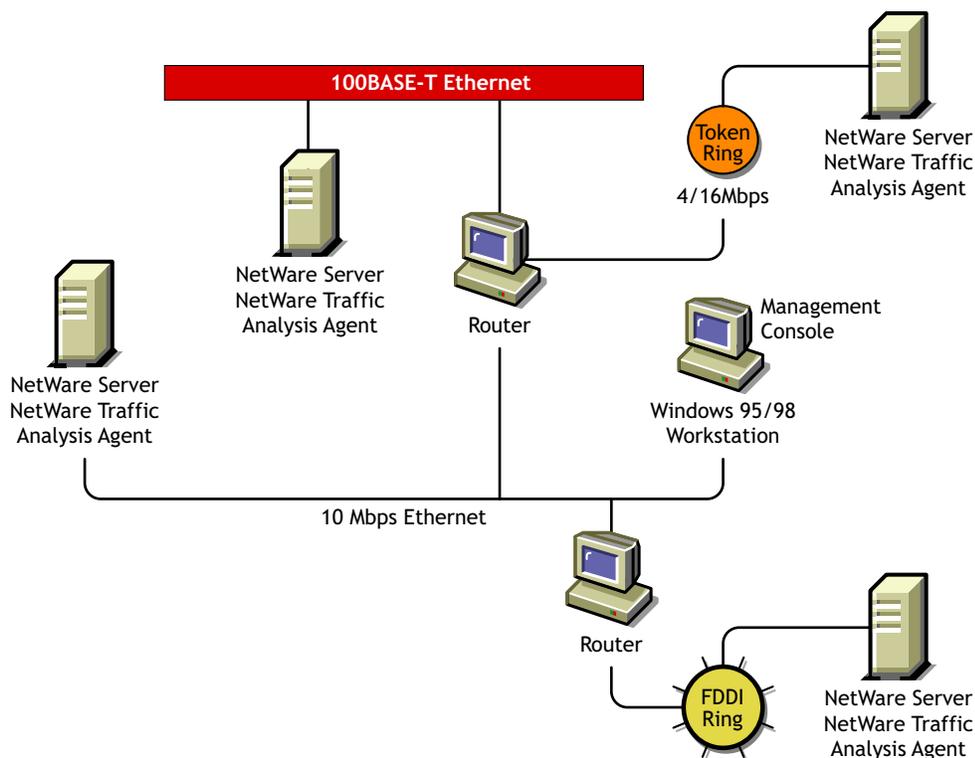
29.7 Using the Traffic Analysis Agent for NetWare

The Traffic Analysis Agent for NetWare (NLA 1.30) runs on a NetWare server. It is a set of NLM programs that enable NetWare to monitor traffic on Ethernet, FDDI, or token ring segments.

The Traffic Analysis Agent for NetWare implements token ring extensions for the RMON MIB ([RFC 1513 \(http://www.isi.edu/in-notes/rfc1513.txt\)](http://www.isi.edu/in-notes/rfc1513.txt)) for token ring media, and a Novell proprietary MIB for FDDI media, in addition to implementing an RMON ([RFC 1757 \(http://www.isi.edu/in-notes/rfc1757.txt\)](http://www.isi.edu/in-notes/rfc1757.txt)) for Ethernet media. The Traffic Analysis Agent for NetWare also implements the first two groups for RMON2 ([RFC 2021 \(http://www.isi.edu/in-notes/rfc2021.txt\)](http://www.isi.edu/in-notes/rfc2021.txt))

Figure 29-8 illustrates a functional view of the Traffic Analysis Agent for NetWare:

Figure 29-8 Traffic analysis agent for NetWare



The following sections provide information about optimizing and using the Traffic Analysis Agent for NetWare:

- ◆ Section 29.7.1, “Planning to Install the Traffic Analysis Agent for NetWare,” on page 1091
- ◆ Section 29.7.2, “Optimizing the Traffic Analysis Agent for NetWare Performance,” on page 1092
- ◆ Section 29.7.3, “Using the Console Utility of the Traffic Analysis Agent for NetWare,” on page 1098

29.7.1 Planning to Install the Traffic Analysis Agent for NetWare

To successfully install the Traffic Analysis Agent for NetWare on a NetWare server, the server must meet the system requirements specified in “[Management and Monitoring Services Installation](#)” in the *Novell ZENworks 7 Server Management Installation Guide*.

You should configure NetWare SNMP parameters as explained in [Chapter 34, “Using SNMP Community Strings,”](#) on page 1159. This will ensure a smooth installation of the Traffic Analysis Agent for NetWare on the server.

NOTE: Although it is not required, it is recommended that you uninstall previous versions of the Traffic Analysis Agent (referred to as the Traffic Analysis Agent in Novell ZENworks Server

Management). If you do not uninstall the previous version of the agent, you must verify that the upgraded NetWare servers run the new Traffic Analysis Agent.

29.7.2 Optimizing the Traffic Analysis Agent for NetWare Performance

The measures described in the following sections can improve the performance of your Traffic Analysis Agent for NetWare servers.

You can configure the Traffic Analysis Agent for NetWare functions described in the following sections by setting the parameters in the `lanz.ncf` file.

- ♦ [“Contents of the LANZ.NCF File” on page 1092](#)
- ♦ [“Modifying the LANZ.NCF File” on page 1095](#)

Contents of the LANZ.NCF File

The `lanz.ncf` file loads all the NLM software required for the Traffic Analysis Agent for NetWare operation. The `lanz.ncf` file resides in the `sys:\zfs_agnt\lanz` directory.

The following example displays the complete text of the default `lanz.ncf` file.

```
#
# Novell NetWare Traffic Analysis Agent
# Version 1.3
#
# - - - - -
# LANZ.NCF: Novell NetWare Traffic Analysis Agent Load File
#
# This NCF file is created by the Novell NetWare Traffic Analysis Agent
install program.
# It is used to load the Novell NetWare Loadable Module files that make
up Novell NetWare
# Traffic Analysis Agent.
# WARNING:  You should not modify this file unless you need to change
one of
# the configuration parameters documented below. Other changes to this
# file are not recommended. Should you damage this file, you must
reinstall
# Novell NetWare Traffic Analysis Agent.
#
# NOTE:      To enable or disable the monitoring of network adapters by
# Novell NetWare Traffic Analysis Agent, use the LANZCON utility as
described in the
# Novell NetWare Traffic Analysis Agent Installation and
Administration guide.
#
# - - - - -
# Load Parameter Descriptions
#
```

```
# load LANZSU debug=1
#
# debug=1    Turns on the LANZ Control screen to see the transactional
# messages from the Novell NetWare Traffic Analysis Agent.
#
# load LANZMEM bound=KB age=HHH
#
# bound=KB   This is the upper limit on memory that can be allocated
# dynamically by the Novell NetWare Traffic Analysis Agent.
#
# Increasing this number allows you to create larger packet
# capture buffers and maintain data for inactive stations
# for a longer period of time.
#
# Decreasing this value reduces the amount of memory that
# can be used by Novell NetWare Traffic Analysis Agent. This leaves
more
# memory for the other server tasks.
#
# Novell NetWare Traffic Analysis Agent automatically purges data for
# inactive stations as the memory boundary is approached.
# This allows Novell NetWare Traffic Analysis Agent to adjust to
#
# the memory that is available to it dynamically.
#
# If the boundary is low, purging occurs frequently, saving
# only data for stations that have been recently active on
# the network. If this happens, a message appears on the
# system console indicating that not enough memory has been
# allocated to Novell NetWare Traffic Analysis Agent.
#
# KB is the memory boundary in kilobytes.
#
# Initial value: Set by the installation program
# based on memory usage
#
# Minimum recommended value:      512
#
# Maximum recommended value:      75% of free server memory
# when NLM files are loaded
#
# Default value:                   If bound=KB is not specified,
# it defaults to 3072.
#
# age=HHH   Novell NetWare Traffic Analysis Agent purges data for
stations that have
```

```

# not been active on the network recently. This parameter
# controls how long data for inactive stations is maintained.
#
# Memory that is used by the station table is not available
# for other uses, such as capturing packets. Reducing the
# AGE value tends to increase the amount of memory
# available for capturing packets.
#
# If you cannot allocate capture buffers that are large,
# you may need to reduce the AGE value.
#
# HHH is the inactivity period, in hours, before station data
# is purged.
#
# Minimum recommended value:      1
#
# Default value:                    If age=HHH is not specified,
# it defaults to 168 (1 week)
#
# load LANZDI level=1
#
# level=1      It indicates that the LANZDI will stop receiving packets
# when CPU utilization gets high.
#
# Default is OFF. LANZDI will continue to receive packets even
# when CPU utilization gets high.
#
# load LANZSM topn=N
#
# topn=N      The number of concurrent sorts of top N nodes that
#
# Novell NetWare Traffic Analysis Agent supports for each network
# adapter.
#
# Recommended value: 4
# Minimum value:    2
# Maximum value:    10
#
# load LANZTR poll = 1
#
# poll=1      Polls token ring source-routed bridges.
#
# load LANZCTL trapreg=1
#
# trapreg=1 Causes SNMP traps to be sent to management consoles
# advertising themselves on the network, as well as stations

```

```

# listed in SYS:\ETC\TRAPTARG.CFG. Omitting this parameter
# or setting it to 0 causes traps to be sent only to those
# stations listed in the SYS:\ETC\TRAPTARG.CFG file.
#
# - - - - -
load gtrend.nlm
load lanzsu.nlm
load lanzmem.nlm bound = 3072 AGE = 168
load lanzlib.nlm
load lanzdi.nlm
load lanzael.nlm
load lanzhis.nlm
load lanzfcb.nlm
load lanzsm.nlm topn = 4
load lanztr.nlm
load lanzfddi.nlm
load lanzctl.nlm trapreg = 1

```

Modifying the LANZ.NCF File

The following sections describe how to modify the parameters of the commands in the `lanz.ncf` file to configure the Traffic Analysis Agent for NetWare functions:

- ◆ “Turning On the LANZ Control Screen” on page 1095
- ◆ “Disabling Packet Capture” on page 1096
- ◆ “Disabling Generation of Duplicate IP Address Alarms” on page 1096
- ◆ “Setting Packet Flow Control” on page 1096
- ◆ “Setting the Upper Limit of Available Memory” on page 1096
- ◆ “Purging Data from Server Memory” on page 1097
- ◆ “Sorting Concurrent Top Stations” on page 1097
- ◆ “Automatically Sending Alarms to the Management Site Server” on page 1098
- ◆ “Polling Source Route Bridges” on page 1098
- ◆ “Activating Changes in the LANZ.NCF File” on page 1098

To make changes in the `lanz.ncf` file and modify the configuration of the Traffic Analysis Agent for NetWare:

- 1 Open the `lanz.ncf` file with a text editor.
- 2 Insert or modify the appropriate parameter as shown and save the file.
- 3 Unload and reload the Traffic Analysis Agent for NetWare, as described in “Activating Changes in the LANZ.NCF File” on page 1098.

Turning On the LANZ Control Screen

The LANZ control screen reports significant events for the Traffic Analysis Agent for NetWare.

To turn on the LANZ control screen, insert the `DEBUG` parameter in the `LOAD LANZSU.NLM` statement, as shown below:

```
LOAD LANZSU.NLM DEBUG=1
```

The default is Off.

Disabling Packet Capture

You might want to disable packet capture to prevent others from observing sensitive data captured in the packets sent on the network segment.

To disable the packet capture, insert a comment mark (#) in the LOAD LANZFCB statement, as shown below:

```
LOAD LANZFCB.NLM
```

You can also control packet capture during high levels of traffic instead of disabling packet capture entirely. For details, see [“Setting Packet Flow Control” on page 1096](#).

Disabling Generation of Duplicate IP Address Alarms

In the DHCP environment, the IP address is released to the DHCP server when a DHCP client is shut down. During the process of releasing the IP address to the DHCP server, the client sends a DHCPRELEASE packet. If this packet does not reach the agent, false duplicate IP address alarms will be generated.

To disable the generation of duplicate IP address alarms, specify zero (0) as the value for the DUPIP parameter, as shown below:

```
LOAD LANZSM DUPIP=0
```

If the DUPIP parameter contains a non-zero value or if the parameter is not specified, duplicate IP address alarms are generated.

Setting Packet Flow Control

The Traffic Analysis Agent for NetWare typically operates in promiscuous mode, receiving all packets on the network. However, if server utilization is high and performance becomes degraded, you can set the LEVEL parameter to 1, which configures the agent to pause when server traffic is high, and then automatically resume operation in promiscuous mode when the traffic level returns to normal.

The default is not to specify the LEVEL parameter at all, which allows continuous operation in promiscuous mode.

To set packet flow control, use the LEVEL parameter setting, as shown below:

```
LOAD LANZDI LEVEL=1
```

Setting the Upper Limit of Available Memory

The BOUND parameter sets the upper limit of available memory that can be allocated dynamically to the Traffic Analysis Agent for NetWare.

The value of the BOUND parameter is measured in kilobytes (KB). The default value is 3072 KB. The minimum recommended value is 512 KB. The maximum recommended value is 75% of the memory that is available after all NLM files are loaded.

You might receive the message “Insufficient memory available for the Traffic Analysis Agent for NetWare” in the following situations:

- ♦ The server has too little memory
- ♦ The server has sufficient memory, but the memory is not available to the Traffic Analysis Agent for NetWare
- ♦ You requested a packet capture buffer that is too large, and the agent granted you less memory than requested

In each case, you should increase the value of the BOUND parameter and add more RAM to your NetWare server.

To change the upper limit of available memory, edit the BOUND parameter, with the appropriate value, as shown below:

```
LOAD LANZMEM BOUND=3072 AGE=168
```

Purging Data from Server Memory

The Traffic Analysis Agent for NetWare holds its data in server memory. You can control the amount of data held in memory by setting the value of the AGE parameter. When data reaches the age specified in the parameter, the data is purged from memory. The AGE parameter is particularly useful on large, bridged networks.

The value of the AGE parameter is measured in hours. The default value is 168, or one week. The minimum recommended value is one hour.

You should lower the AGE parameter if you receive the message “Insufficient memory available for the Traffic Analysis Agent for NetWare” and you have allocated sufficient memory for the agent.

Having insufficient memory is not harmful to the agent or the server. The Traffic Analysis Agent for NetWare can run indefinitely, even when the memory allocated to it is not sufficient.

To modify the amount of data held in server memory, change the value of the AGE parameter, as shown below:

```
LOAD LANZMEM BOUND=3072 AGE=168
```

Sorting Concurrent Top Stations

The Traffic Analysis Agent for NetWare sorts stations whenever the top eight graphs on the Segment Dashboard view, the Stations view, or both are displayed by Novell ConsoleOne. The sorts are independent of each other and can be computed on the basis of different statistics.

Because each of the sort computations uses server CPU cycles, you should limit the number of concurrent computations.

To set the number of concurrent sort computations per network adapter, set the TOPN parameter, as shown below:

```
LOAD LANZSM TOPN=n
```

The default value is 4. The minimum value is 2. The maximum value is 10.

Automatically Sending Alarms to the Management Site Server

The Traffic Analysis Agent for NetWare can automatically send SNMP alarms (sometimes referred to as SNMP traps) to the Management Site Server or other nodes on the network in the following configurations:

- ◆ The Traffic Analysis Agent for NetWare receives the SAP packets sent by the Management Site Server.
- ◆ The Management Site Server or other node is listed in the server's `traptarg.cfg` file. This file can be edited to add other trap targets.

The `traptarg.cfg` file is stored in the `sys:\etc` directory. The file provides instructions for its use. You can edit the file with any ASCII text editor.

To enable alarms to be sent automatically, add the TRAPREG parameter setting, as shown below:

```
LOAD LANZCTL TRAPREG=1
```

The default is 1. If you omit the TRAPREG parameter or set its value to zero (0), the agent sends alarms only to management consoles listed in the `traptarg.cfg` file.

Polling Source Route Bridges

To control source route bridge polling on token ring networks, use the POLL parameter, as shown below:

```
LOAD LANZTR POLL=1
```

1 = On and 0 = Off.

Setting the POLL parameter to 1 polls source routed bridges once every second. You cannot change the polling rate. The default is On.

To turn off this function, set the POLL parameter to zero (0), as shown below:

```
LOAD LANZTR POLL=0
```

The default is to omit the POLL parameter. Also, the LOAD LANZTR statement is commented out on systems that do not have a token ring adapter installed.

Activating Changes in the LANZ.NCF File

To activate the changes you make in the `lanz.ncf` file:

- 1 Save the LANZNCF file.
- 2 Enter ULANZ at the server prompt to unload the agent.
- 3 Enter LANZ to reload the agent.

29.7.3 Using the Console Utility of the Traffic Analysis Agent for NetWare

The Traffic Analysis Agent for NetWare 1.3 provides a console utility (`lanzcon.nlm`) that performs the following three tasks:

- ◆ Enables or disables network monitoring by the selected network adapters
- ◆ Provides a source of detailed troubleshooting information

- ◆ Resolves a residual entry (for example, a Host TopN entry created by a management console that terminated unexpectedly)

When you install the Traffic Analysis Agent for NetWare, `lanzcon.nlm` is installed automatically in the `sys:\zfs_agnt\lanz` directory.

The following topics are discussed in greater detail in this section:

- ◆ “Loading the Console Utility of the Traffic Analysis Agent for NetWare” on page 1099
- ◆ “Enabling or Disabling Network Adapter Monitoring” on page 1099
- ◆ “Viewing Network Adapter Information” on page 1100
- ◆ “Viewing the Agent Item Status” on page 1100
- ◆ “Accessing Detailed Information About Each Item” on page 1101
- ◆ “Migrating Trend Files” on page 1104

Loading the Console Utility of the Traffic Analysis Agent for NetWare

To use `lanzcon.nlm`, enter the following command at the NetWare console prompt:

```
LOAD LANZCON CONTROLCOMMUNITY = <control community string>
```

IMPORTANT: If LANZCON is launched without any command line argument, then the default control community string is PUBLIC.

`lanzcon.nlm` is loaded and displays a list of network adapters, along with summary information about the network adapters currently installed on the server.

The following information is displayed for each network adapter:

- ◆ **Number (#):** The network adapter entry number in the network interface table.
- ◆ **Description:** A brief description of the network adapter.
- ◆ **Media Type:** The type of network connected to the network adapter: Ethernet, FDDI, or token ring.
- ◆ **Adapter Address:** The physical address of the network adapter.

Enabling or Disabling Network Adapter Monitoring

To enable or disable monitoring of a selected network adapter:

- 1 From the Network Adapters screen, select the appropriate adapter then press F3.
 - ◆ If the selected adapter is currently monitoring an Ethernet or token ring network, the console displays the Adapter Is Monitoring screen.
 - ◆ If the selected adapter is not monitoring an Ethernet or token ring network, the console displays the Adapter Is Not Monitoring screen.
- 2 Select *Yes* or *No* to enable or disable monitoring.

If you disable monitoring, all LAN analysis data for the selected adapter is deleted.

Using LANZCON, an FDDI adapter cannot be disabled. To disable an FDDI adapter:

- 1 Unload LANZCON, if loaded.

- 2 Unload LANZ, if loaded.
- 3 Open `lanz.ncf` from `sys:\zfs_agn\lanz` directory for editing.
- 4 Comment the statement `LOAD lanzfddi.nlm` by entering the # symbol at the beginning of this statement.
- 5 Save `lanz.ncf` and exit.
- 6 Reload LANZ.

Viewing Network Adapter Information

To bring up detailed information for network adapter items:

- 1 From the Network Adapters screen, select an adapter then press Enter.
- 2 From the Select Information to View screen, select Show Adapter Items.

The LANZCON utility displays the Network Adapter Items screen that lists all the items related to the selected network adapter.

The screen for a token ring adapter includes the information from the Novell Token Ring RMON MIB. For details, see [“Viewing the Agent Item Status” on page 1100](#).

To return to the Select Information to View menu, press Esc.

The following information is provided for the selected adapter:

- ♦ **Item:** The types of items that are currently being monitored by the selected adapter. The Network Adapter Items screen shows a set of typical items consisting of token ring, Statistics, History, Host, Matrix, and Host TopN. The Traffic Analysis Agent for NetWare monitors these items by default. In the Network Adapter Items screen, the Host TopN item, indicating the list of the busiest nodes, has been added by a user. You can add other items to this display in Novell ConsoleOne, depending on your configuration.

You can select any item to view more information about each topic. To view the values for the selected item, select the desired item then press Enter. Refer to the following sections for more examples of the screens.

- ♦ **Index:** The entry number of the displayed item in the list of all the items of the same type. The related tables are identified by this index.
- ♦ **Description:** A textual description of the entry. This column indicates the software entity or user that created the item. The items automatically monitored by the Traffic Analysis Agent for NetWare are indicated by the monitor.

For a token ring network entry, this column shows the media speed and the local ring number.

Viewing the Agent Item Status

When you click the *Select Information to View* menu > *Show Agent Items*, LANZCON displays all the items for each network adapter being monitored by the Traffic Analysis Agent for NetWare.

To view the agent item status for the selected agent:

- 1 From the Network Adapters screen, select an adapter then press Enter.
- 2 From the Select Information to View screen, select Show Agent Items.

The All Novell NetWare Traffic Analysis Agent Items screen shows all the items related to the agent monitoring the segment. For example, if you are using multiple adapters to monitor multiple network segments, the screen lists all the items being monitored by the agent.

To delete any entry (except the token ring network entry), select the entry then click *Delete*, and then click *Yes*.

To return to the Network Adapter Items screen, press Esc.

The following information is provided for the agent:

- ◆ **Item:** The types of items available. The All Novell NetWare Traffic Analysis Agent Items screen shows a set of typical items consisting of Statistics, History, Host, Matrix, and Host TopN. Additional items can be displayed, depending on your configuration.

You can select any item for more information about each topic. To view the values for an item, select the desired item then press Enter. See the following sections for more examples of the screens.

- ◆ **Index:** The entry number of the displayed item in the list of all items of the same type. The related tables are identified by this index.
- ◆ **Description:** A textual description of the entry. This column indicates the software entity or user that created the item table. The items automatically monitored by the Traffic Analysis Agent for NetWare are indicated by the monitor.

For a token ring network entry, this column shows the media speed and the local ring number.

Accessing Detailed Information About Each Item

This section describes the major categories of information available for both the selected network adapter and the Traffic Analysis Agent for NetWare. The following topics are covered:

- ◆ [“Viewing the Token Ring RMON MIB Information” on page 1101](#)
- ◆ [“Viewing the FDDI Ring RMON MIB Information” on page 1101](#)
- ◆ [“Viewing Statistics Information” on page 1102](#)
- ◆ [“Viewing History Information” on page 1102](#)
- ◆ [“Viewing Host Information” on page 1102](#)
- ◆ [“Viewing Matrix Information” on page 1103](#)

Viewing the Token Ring RMON MIB Information

To view the Token Ring RMON MIB information:

- 1 From the Network Adapter Items screen, select the token ring item, then press Enter.
- 2 From the Select Information to View screen, select *Show Adapter Items*, then press Enter.
- 3 Press Esc to exit this screen.

Viewing the FDDI Ring RMON MIB Information

To view the FDDI ring RMON MIB information:

- 1 From the Network Adapter Items screen, select the *FDDI Ring* item, then press Enter.
- 2 From the Select Information to View screen, select *Show Adapter Items*, then press Enter.

Viewing Statistics Information

The statistics information presents the basic statistics for each monitored adapter per segment.

To view the statistics information:

- 1 From the Network Adapter Items screen, select *Statistics*.
- 2 Press Enter.

For an Ethernet network entry, the LANZCON utility displays the Statistics Information screen.

This screen displays the statistical values of the selected network adapter. The display is updated periodically with the latest values for each field.

- 3 To exit this screen, press Esc.

Viewing History Information

The history information defines sampling functions for the networks that are being monitored. The History Control table defines a set of samples at a particular sampling interval for a particular network adapter.

To view the history information:

- 1 From the Network Adapter Items screen, select *History*, then press Enter.
- 2 To exit this screen, press Esc.

The field descriptions are as follows:

- ♦ **Index:** An integer that uniquely identifies a row in the History Control table.
- ♦ **Data Source:** Identifies the network adapter and the Ethernet, FDDI, or token ring segment that is the source of the data for entries defined by this object.
- ♦ **Buckets Requested:** The requested number of discrete sampling intervals over which data will be saved in the portion of the media-specific table associated with this entry.
- ♦ **Buckets Granted:** The actual number of discrete sampling intervals over which data will be saved.
- ♦ **Interval:** The interval, in seconds, over which data is sampled for each bucket. The interval can be set to any number between 1 and 3,600 (one hour). The default interval for past hour is 30 seconds per sample, and the default interval for past day is 30 minutes (or 1,800 seconds) per sample.

The sampling scheme is determined by the buckets granted and the control interval.

- ♦ **Owner:** The entity that created the item. "Monitor" indicates that the item was created by the Traffic Analysis Agent for NetWare.
- ♦ **Status:** A status of Valid indicates that the agent is operating normally under the instructions given by the table.

Viewing Host Information

The host group gathers statistics about specific hosts or nodes on the LAN. The Traffic Analysis Agent for NetWare learns of new nodes on the LAN by observing the source and destination MAC addresses in good packets. For each node known to the agent, a set of statistics is maintained.

To view the host (node) information:

1 From the Network Adapter Items screen, select *Host* then press Enter.

The host group consists of three tables: two data tables and one control table. The two data tables are *hostTable* and *hostTimeTable*. The control table, *hostControlTable*, includes the following objects, which correspond to the fields displayed in the Host Information screen:

- ◆ **Index:** An integer that uniquely identifies a row in the *hostControlTable*. Each row in the control table refers to a unique network adapter, and thus, a unique segment.
- ◆ **Data Source:** Identifies the network adapter and the Ethernet, FDDI, or token ring segment that is the source of the data for the entries defined by this object.
- ◆ **Table Size:** The number of rows in the *hostTable* associated with this row.
- ◆ **Last Delete Time:** The value of the *sysUpTime* MIB object that corresponds to the last time an entry was deleted from the portion of the *hostTable* associated with this row. The value is zero (0) if no deletions occurred.
- ◆ **Owner:** Indicates the entity or user that created the item. “Monitor” indicates that the item was created by the Traffic Analysis Agent for NetWare.
- ◆ **Status:** A status of Valid indicates that the agent is operating normally under the instructions given by the table.

Viewing Matrix Information

The matrix group records information about the conversations between pairs of nodes on a network segment. The information is stored in the form of a matrix. This method of organization is useful to retrieve specific pairings of traffic information, such as finding out which nodes are making the most use of a server.

To view the matrix information:

1 From the Network Adapter Items screen, select *Matrix*, then press Enter.

The matrix group consists of three tables: two data tables and one control table. The data tables are *matrixSDTable* and *matrixDSTable*. The control table, *matrixControlTable*, includes the following objects, which correspond to the fields displayed in the Matrix Information screen:

- ◆ **Index:** An integer that uniquely identifies a row in the *matrixControlTable*. Each row in the control table defines a function that discovers conversations on a particular network and places statistics about them in the two data tables.
- ◆ **Data Source:** Identifies the network adapter, and the Ethernet, FDDI, or token ring segment that are the source of the data for the entries defined by this object.
- ◆ **Table Size:** The number of rows in the *matrixTable* associated with this row.
- ◆ **Last Delete Time:** The value of the *sysUpTime* object that corresponds to the last time an entry was deleted from the portion of the *matrixTable* associated with this row. The value is zero (0) if no deletions occurred.
- ◆ **Owner:** Indicates the entity or user that created the item. “Monitor” indicates that the item was created by the Traffic Analysis Agent for NetWare.
- ◆ **Status:** A status of Valid indicates that the agent is operating normally under the instructions given by the table.

Migrating Trend Files

In Novell ConsoleOne, you can view trends of traffic patterns on the monitored Ethernet, FDDI, and token ring segments. You can use the trend data to analyze traffic on the segment. For details, see [“Analyzing Trend Data for a Segment” on page 1050](#).

Earlier versions of the Traffic Analysis Agent for NetWare (1.20 and 1.21) collected trend data that was sampled every one minute. The Traffic Analysis Agent for NetWare 1.30 that ships with Novell ZENworks Server Management collects trend data that are sampled every one minute, one hour, and one day. This functionality of version 1.30 of the Traffic Analysis Agent for NetWare ensures minimal communication between the agent and Novell ConsoleOne, to reduce network traffic.

You can use the migrating tool (`gtrend.exe`) to convert the trend data collected by earlier versions of the Traffic Analysis Agent for NetWare to trend data that can be used by version 1.30 of the Traffic Analysis Agent for NetWare and Novell ConsoleOne.

To migrate trend files collected by versions 1.20 or 1.21 of the Traffic Analysis Agent for NetWare:

- 1 Copy `gtrend.exe` from the Installation CD to a TEMP folder on a 32-bit Windows machine.
- 2 Copy the trend data files collected by earlier versions of the Traffic Analysis Agent for NetWare to the TEMP folder.
- 3 Run `gtrend.exe`.

This will migrate the existing one-minute trend files to the corresponding one-hour and one-day trend files that can be used by version 1.30 of the Traffic Analysis Agent for NetWare.

- 4 Copy the migrated trend files to the `sys:\gtrend\` folder on the NetWare server and run the version 1.30 of the Traffic Analysis Agent for NetWare on the same server.

NOTE: The migration tool will not migrate older token ring trend data collected by version 1.20 or 1.21 of the Traffic Analysis Agent for NetWare because the older agents implemented a proprietary Token Ring MIB that enabled the agent to collect trend data sampled every one minute. Version 1.3 of the Traffic Analysis Agent for NetWare implements the standard Token Ring MIB that supports historical trends (one minute, one hour and one day).

29.8 Using the Traffic Analysis Agent for Windows

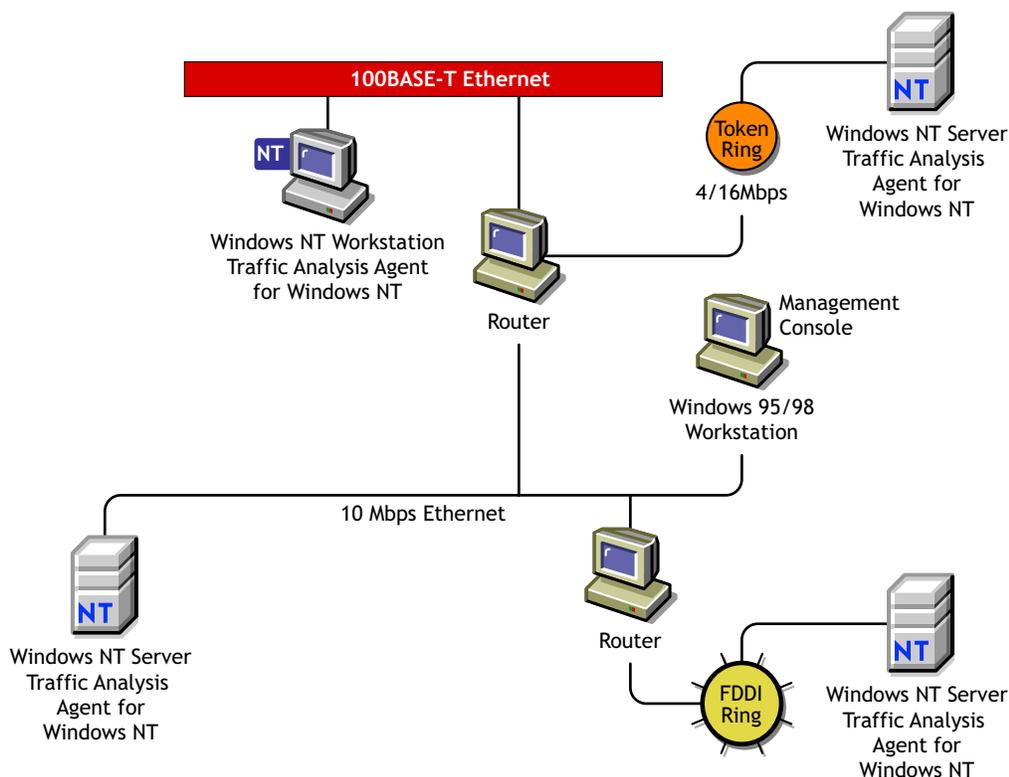
The Traffic Analysis Agent (version 1.30) for Windows runs on a Windows server or on a Windows workstation. The Traffic Analysis Agent for Windows monitors traffic on Ethernet, FDDI, or token ring segments.

The Traffic Analysis Agent for Windows is an RMON agent that implements functionality defined by the RMON MIB. It implements token ring extensions for RMON ([RFC 1513 \(http://www.isi.edu/in-notes/rfc1513.txt\)](http://www.isi.edu/in-notes/rfc1513.txt)) for token ring media, and a Novell proprietary MIB for FDDI media, in addition to implementing an RMON ([RFC 1757 \(http://www.isi.edu/in-notes/rfc1757.txt\)](http://www.isi.edu/in-notes/rfc1757.txt)) for Ethernet media.

The agent collects information about activity on your network and makes it available to Novell ConsoleOne via SNMP. The Traffic Analysis Agent for Windows also implements the first two groups of RMON2 ([RFC 2021 \(http://www.isi.edu/in-notes/rfc2021.txt\)](http://www.isi.edu/in-notes/rfc2021.txt)).

Figure 29-9 illustrates a functional view of the Traffic Analysis Agent for Windows:

Figure 29-9 Traffic Analysis Agent for Windows



The following sections provide information about optimizing and using the Traffic Analysis Agent for Windows:

- ◆ [Section 29.8.1, “Changes Made During Installation,”](#) on page 1105
- ◆ [Section 29.8.2, “Planning to Install the Traffic Analysis Agent for Windows,”](#) on page 1106
- ◆ [Section 29.8.3, “Optimizing the Traffic Analysis Agent for Windows,”](#) on page 1109
- ◆ [Section 29.8.4, “Using LANZCON,”](#) on page 1112

29.8.1 Changes Made During Installation

When you install the Traffic Analysis Agent for Windows, the files listed in [Table 29-34](#) are copied to Windows:

Table 29-34 Files are copied to Windows during the Traffic Analysis Agent for Windows installation

| Filename | Location | Description |
|--------------|-------------------------|------------------------------|
| lanzndis.sys | \winnt\system32\drivers | Kernel mode driver interface |
| lanzctl.dll | \winnt\system32 | Control module |
| lanzmem.dll | \winnt\system32 | Memory manager module |
| lanzlib.dll | \winnt\system32 | Library module |

| Filename | Location | Description |
|--------------|-----------------------|--|
| lanzdi.dll | \winnt\system32 | User mode driver interface |
| lanzsm.dll | \winnt\system32 | Monitor module |
| lanzhis.dll | \winnt\system32 | History module |
| lanzael.dll | \winnt\system32 | Alarm, event, and log module |
| lanzfcbl.dll | \winnt\system32 | Filter capture, buffer module |
| lanztr.dll | \winnt\system32 | Token ring manager module |
| lanzfdi.dll | \winnt\system32 | FDDI manager module |
| gtrend.dll | \winnt\system32 | Trend module |
| lanzcon.exe | \lanznt | Agent console application |
| lanzcon.chm | \lanznt | Agent console help |
| gtrend.exe | \zfs_agnt\lanzcon | Tool for migration of trend data from the older agent. |
| mgmtapi.dll | \zfs_agnt\lanzcon | SNMP application file |
| msvcpl50.dll | \zfs_agnt\lanzcon | MFC APIs required for LANZCON |
| lanzctl.dll | \zfs_agnt\lanzcon | Required for LANZCON |
| msflxgrd.ocx | %systemroot%\system32 | Enables ActiveX* Controls in LANZCON |

IMPORTANT: The default directory location for the LANZCON application is `zfs_agnt\lanzcon`. You can change the location of LANZCON during installation.

29.8.2 Planning to Install the Traffic Analysis Agent for Windows

The Traffic Analysis Agent for Windows requires configuration of the Windows SNMP service before installing the agent.

- ◆ [“Installing and Configuring the Windows SNMP Service” on page 1106](#)
- ◆ [“Installing and Configuring the Windows 2000 SNMP Service” on page 1107](#)

Installing and Configuring the Windows SNMP Service

Before installing the Novell ZENworks Server Management agent, you must install and configure the Windows SNMP service. This is required to enable communication with the management server.

To install and configure SNMP on Windows:

- 1 Install the SNMP service.
 - 1a In the Control Panel, select *Network > Services > click Add*.
 - 1b Select *SNMP Service* from the Select Network Service dialog box.
 - 1c Click *OK*.

To install and configure SNMP on Windows 2000:

- 1** Install the SNMP service.
 - 1a** In the Control Panel, select *Administrative Tools*, then click *Configure Your Server*.
 - 1b** In the *Application Server* option, select *Terminal Services*.
 - 1c** Click *Start*.
 - 1d** In the Windows Components Wizard, double-click *Management and Monitoring Tools*.
 - 1e** Select *Simple Network Management Protocol*.
 - 1f** Click *OK*.
 - 1g** Click *Next*.

SNMP is started automatically after installation.
- 2** Configure the SNMP Trap service to start automatically.
 - 2a** In the Control Panel, select *Administrative Tools > Services*.
 - 2b** Click *SNMP Trap Service > Startup*.
 - 2c** In the *Startup Type* options, select *Automatic*.
- 3** Specify the trap community name and trap destination address so that the agent sends traps to the management server.
 - 3a** In the Control Panel, select *Administrative Tools > Services*.
 - 3b** Double-click *SNMP Service*.
 - 3c** Click *Properties*.
 - 3d** Click the *Traps* tab.
 - 3e** Select a name from the *Community Names* box, then click *Add*.

The *Add* button is disabled if there are no Community Names available.
 - 3f** If the public community name is not present, enter `public`.
 - 3g** Click *Add*.
 - 3h** Use the *Trap Destinations* box to add other DNS names and IP addresses in addition to the loopback IP address for the workstations or servers that should receive traps.
 - 3i** Click *OK*.
- 4** Set the SNMP security options trap community name so that SNMP packets from any host are accepted by the agent.
 - 4a** In the Control Panel, select *Administrative Tools > Services*.
 - 4b** Double-click *SNMP Service*.
 - 4c** Click *Properties*.
 - 4d** Click the *Security* tab.
 - 4e** In the *Accepted Community Names* box, click *Add*.
 - 4f** Select a name from the *Community Name* box.

The *Accepted Community Names* list displays the community names from which Windows 2000 will accept requests.
 - 4g** Click *Add*.
 - 4h** Select *Accept SNMP Packets from Any Host > click OK*.

IMPORTANT: After installing the SNMP services, you should re-install the service packs again.

29.8.3 Optimizing the Traffic Analysis Agent for Windows

The Traffic Analysis Agent for Windows parameters are configured for optimal performance on Windows. You can optimize the performance of the agent to suit your networking environment.

This section explains how to optimize the agent and monitor the functionality Traffic Analysis Agent for Windows using the agent console (LANZCON) for Windows. For details, see [“Using LANZCON” on page 1112](#).

The following sections explain the Traffic Analysis Agent for Windows configuration options:

- ◆ [“Configuring the Traffic Analysis Agent for Windows” on page 1109](#)
- ◆ [“Configuring the Modules of the Traffic Analysis Agent for Windows” on page 1109](#)
- ◆ [“Configuring the Parameters of the Traffic Analysis Agent for Windows” on page 1110](#)
- ◆ [“Automatically Loading the Agent with the SNMP Service” on page 1111](#)

Configuring the Traffic Analysis Agent for Windows

The Traffic Analysis Agent for Windows provides default values for modules and parameters. You can change the default values to optimize the performance of the Traffic Analysis Agent for Windows.

You can configure the following modules of the Traffic Analysis Agent for Windows:

- ◆ Packet Capture
- ◆ Station Monitor
- ◆ Token Ring Manager
- ◆ FDDI Manager

For details, see [“Configuring the Modules of the Traffic Analysis Agent for Windows” on page 1109](#).

You can configure the following parameters of the Traffic Analysis Agent for Windows:

- ◆ Memory Bound
- ◆ Memory Age
- ◆ Top N Station
- ◆ Generate Duplicate IP Address Alarms
- ◆ Trend Files Location

For details, see [“Configuring the Parameters of the Traffic Analysis Agent for Windows” on page 1110](#).

Configuring the Modules of the Traffic Analysis Agent for Windows

By default, all agent modules are enabled to load. You can choose to disable the modules.

To disable the modules of the Traffic Analysis Agent for Windows:

- 1 From the LANZCON main menu, click *Configure > Traffic Analysis Agent Modules > Disable*.
- 2 Deselect the module you want the agent to monitor.
- 3 Click *OK*.

Configuring the Parameters of the Traffic Analysis Agent for Windows

The Traffic Analysis Agent for Windows modules are loaded with default parameters. You can modify the parameters to optimize the performance of the agent.

Table 29-35 describes the parameters of the Memory Manager module:

Table 29-35 *Parameters of the Memory Manager module*

| Parameter | Default Value | Range | Description |
|--------------|---------------|--------------------|---|
| Memory Bound | 4 MB | 1 MB - 10 MB | Sets the upper limit of available memory that can be allocated dynamically to the Traffic Analysis Agent for Windows. |
| Memory Age | 168 hours | 1 hour - 720 hours | Controls the duration for which the Traffic Analysis Agent for Windows stores data in memory. When the duration setting is reached, existing data is purged from memory. |

To modify the Memory Bound parameter:

- 1 From the LANZCON main menu, click *Configure > Traffic Analysis Agent Parameters*.
- 2 Click the *Memory Manager* tab.
- 3 Move the *Memory Bound* slider to the point you want to set as the memory bound value.

To modify the Memory Age parameter:

- 1 From the LANZCON main menu, click *Configure > Traffic Analysis Agent Parameters*.
- 2 Click the *Memory Manager* tab.
- 3 Move the Memory Age slider to the point you want to set as the memory age value.

IMPORTANT: Restart the Traffic Analysis Agent for Windows to ensure that the agent utilizes the changed parameter values. For details, see “**Management and Monitoring Services Installation**” in the *Novell ZENworks 7 Server Management Installation Guide*.

Table 29-36 describes the parameters of the Station Monitor module:

Table 29-36 Parameters of the Station Monitor module

| Parameter | Default Value | Range | Description |
|--------------------------------------|---------------|----------------|---|
| TopN Station | 4 reports | 2 - 10 reports | Controls the number of TopN reports the agent can generate. |
| Generate Duplicate IP Address Alarms | On | - | Controls the generation of duplicate IP address alarms. |

To specify the number of TopN reports you want the agent to generate:

- 1 From the *LANZCON* main menu, click *Configure > Traffic Analysis Agent Parameters*.
- 2 Click the *Station Monitor* tab.
- 3 Select the number of TopN reports.

To stop generation of duplicate IP address alarms:

- 1 From the *LANZCON* main menu, click *Configure > Traffic Analysis Agent Parameters*.
- 2 Click the *Station Monitor* tab.
- 3 Deselect the *Generate Duplicate IP Address Alarms* check box.

Table 29-37 describes the Network Trend parameter:

Table 29-37 Network Trend parameter

| Parameter | Default Path | Description |
|----------------------|---------------------------|---|
| Trend Files Location | <i>system root\GTREND</i> | Specifies the directory path and location where trend files (*.GT) are created and updated. |

IMPORTANT: If you delete the *.GT file, all the previous trend information will be lost.

To specify a path to a location for storing trend data:

- 1 From the *LANZCON* main menu, click *Configure > Traffic Analysis Agent Parameters*.
- 2 Click the *Network Trends* tab.
- 3 Enter or browse to select the directory path to the location where you want the Traffic Analysis Agent for Windows to store trend data.

Automatically Loading the Agent with the SNMP Service

The Traffic Analysis Agent depends on the Microsoft* SNMP service on Windows. When SNMP starts, it loads agent DLLs in its address space. After the agent is installed, it will be always loaded by the SNMP service, by default, whenever the service starts.

You can enable or disable loading of the agent DLLs with SNMP by checking the desired options in the Novell Traffic Analysis Agent Loading with SNMP dialog box. If you disable the agent, the SNMP service will start normally but the Traffic Analysis Agent will not work. The Traffic Analysis Agent will neither capture packets by placing the NIC cards into the promiscuous mode nor will respond to SNMP requests.

29.8.4 Using LANZCON

This section explains how you can use the LANZCON utility to configure and diagnose the Traffic Analysis Agent for Windows.

LANZCON for Windows is a graphical user interface provided by the Traffic Analysis Agent for Windows to configure the agent modules and parameters and to diagnose the agent. You can use LANZCON to obtain information about network segments monitored by the agent to help you troubleshoot problems.

To open the LANZCON utility, do one of the following:

- ◆ From the Windows *Programs* menu, click *Traffic Analysis Agent for Windows > LANZCON*.
- ◆ Double-click the *LANZCON* icon  on your desktop.

To use LANZCON utility with different SNMP community strings, do the following:

- 1 Go to the `lanzcon` directory.
- 2 Enter the following at the command prompt:

```
LANZCON <community name>
```

IMPORTANT: If you launch LANZCON without using the command line argument, the default community string is PUBLIC.

You can perform the following tasks with LANZCON:

- ◆ [“Viewing Network Adapters” on page 1112](#)
- ◆ [“Enabling or Disabling Network Adapter Monitoring” on page 1113](#)
- ◆ [“Viewing the Agent Log” on page 1113](#)
- ◆ [“Viewing the Agent Status” on page 1113](#)
- ◆ [“Viewing RMON Tables” on page 1114](#)
- ◆ [“Viewing SNMP Traps” on page 1114](#)

Viewing Network Adapters

On loading LANZCON, you will see the Network Adapters window. The Network Adapters window displays information about monitored adapters in two panes.

[Table 29-38](#) describes the two panes in the Network Adapters window:

Table 29-38 *Panes in the Network Adapters window*

| Pane | Displays | Description |
|-----------|-------------------|--|
| Left pane | Adapter Tree view | Displays a list of network adapters discovered by the Traffic Analysis Agent for Windows. The default view displays a collapsed tree. You can expand each network adapter in the tree to view the list of RMON tables for the selected adapter. |

| Pane | Displays | Description |
|------------|------------|--|
| Right pane | Table view | <p>Displays details about the object you select in the left pane.</p> <p>If you select an adapter in the left pane, interface table (RFC 1213 (http://www.isi.edu/in-notes/rfc1213.txt)) details such as media type, MAC address, and description of the selected adapter are displayed in the right pane.</p> <p>If you select an RMON table in the left pane, table data is displayed in the right pane.</p> |

Enabling or Disabling Network Adapter Monitoring

The Traffic Analysis Agent for Windows collects information about monitored adapters and displays it in the right pane of the Network Adapters window.

By default, adapter monitoring is enabled. LANZCON lets you disable adapter monitoring. If you disable adapter monitoring, the Traffic Analysis Agent for Windows stops collecting data for the adapter and the RMON tables for the adapter will be deleted.

IMPORTANT: You cannot disable monitoring FDDI adapters through LANZCON.

To enable adapter monitoring:

- 1 Select an adapter in the left pane of the Network Adapters window.
- 2 Click *View > NetWork Adapters > Enable*.

To disable adapter monitoring:

- 1 Select an adapter in the left pane of the Network Adapters window.
- 2 Click *View > NetWork Adapters > Disable*.

Viewing the Agent Log

The Traffic Analysis Agent for Windows logs significant events and error messages that occurred during a session.

To view the agent log:

- 1 From the *LANZCON* main menu, click *View > Agent Log*.

Viewing the Agent Status

You can view the status of the agent from the Traffic Analysis Agent Status window. The agent status window indicates whether the agent modules are loaded or not loaded.

To view the agent status:

- 1 From the *LANZCON* main menu, click *View > Agent Status*.

Viewing RMON Tables

RMON tables are listed under each network adapter. You can view the RMON tables by selecting a table in the left pane of the Network Adapters window. RMON table data is displayed in the right pane.

The Network Adapter tree displays the following RMON tables:

- ◆ Statistics
- ◆ History Control
- ◆ History Data
- ◆ Host Control
- ◆ Host Entry
- ◆ Host TopN Control
- ◆ Host TopN Entry
- ◆ Matrix Control
- ◆ Matrix SD Entry
- ◆ Filter, Channel, and Buffer

The Alarm Information tree displays the following RMON tables:

- ◆ Alarm
- ◆ Event
- ◆ Log

Viewing SNMP Traps

The Traffic Analysis Agent for Windows monitors network segments and sends traps to the management server. Novell ConsoleOne displays the alarm when it receives the trap from the management server.

Trap information is displayed in the SNMP Traps window. For each trap, [Table 29-39](#) shows trap data that can be obtained.

Table 29-39 Trap information displayed in the SNMP Traps window

| Statistic | Explanation |
|--------------|--|
| Receive Time | Displays the time when the trap occurred |
| Trap Summary | Displays a description of the trap |

IMPORTANT: LANZCON will receive trap notifications if you have ensured that Windows SNMP has been configured to send traps to a loopback trap destination address. For details, see [“Planning to Install the Traffic Analysis Agent for Windows” on page 1106](#).

To view SNMP traps from *LANZCON* main menu, click *View > SNMP Traps*.

Customizing the Agent Configuration

30

The Novell® ZENworks® Server Management with SNMP agents run on Novell NetWare® and Windows* servers in your network. The agents monitor servers, collecting historical data and dynamic data in response to requests from Novell ConsoleOne®. An administrator at the Novell ZENworks Server Management Novell ConsoleOne can request data simply by clicking a representative icon for any device, operating system, or service discovered on a server.

After the Management Agent for NetWare and the Management Agent for Windows have been installed on your network NetWare and Windows servers, they are ready to operate with the default settings. In most cases, this configuration is sufficient; however, you can customize the agent settings to enhance management functionality.

This appendix contains the following sections:

- ◆ [Section 30.1, “Agent Files,” on page 1115](#)
- ◆ [Section 30.2, “Customizing the Management Agent for NetWare,” on page 1118](#)
- ◆ [Section 30.3, “Customizing the Management Agent for Windows Server,” on page 1120](#)
- ◆ [Section 30.4, “Third-Party Agent Configuration,” on page 1122](#)
- ◆ [Section 30.5, “Advanced Trending Agent,” on page 1123](#)
- ◆ [Section 30.6, “Management and Monitoring Services for Linux,” on page 1129](#)

30.1 Agent Files

The following sections describe the agent files that are installed on each managed server:

- ◆ [Section 30.1.1, “Management Agent for NetWare Files,” on page 1115](#)
- ◆ [Section 30.1.2, “Management Agent for Windows Server Files,” on page 1117](#)

30.1.1 Management Agent for NetWare Files

Table 30-1 describes the Management Agent for Novell NetWare NLM™ files installed on a NetWare server:

Table 30-1 List of Management Agent for Novell NetWare NLM files installed on a NetWare server

| Management Agent for NetWare NLM Files | Description |
|--|--|
| <code>servinst.nlm</code> | Implements the NetWare server MIB (<code>nwserver.mib</code>). |
| <code>hostmib.nlm</code> | Implements the standard Host Resources MIB [RFC 1514] and Novell® extensions to that MIB (<code>nwhostx.mib</code>). |

| Management Agent for NetWare NLM Files | Description |
|--|---|
| <code>ntrend.nlm</code> | Implements the Threshold and Trend MIB (<code>nwtrend.mib</code>). When loaded, <code>ntrend.nlm</code> sets trends and thresholds for each monitored attribute according to the server's configuration. The <code>ntrend.ini</code> file contains configuration parameters for <code>ntrend.nlm</code> . |
| <code>nwtrap.nlm</code> | Implements the NetWare Server Trap MIB (<code>nwalarm.mib</code>). The <code>nwtrap.cfg</code> file contains configuration parameters for <code>nwtrap.nlm</code> . |
| <code>findnms.nlm</code> | Used by NetWare servers running the Management Agent for NetWare. Employ <code>findnms.nlm</code> to listen for SNMP Management console advertising themselves using the Service Advertising Protocol (SAP) number 0x026a. <code>Findnms.nlm</code> then adds the Internetwork Packet Exchange™ (IPX™) address of each Novell ConsoleOne discovered to the list of stations that receive traps. |
| <code>ndstrap.nlm</code> | Implements the <code>ndstrap.mib</code> to capture and forward Novell eDirectory events to SNMP Management console. |
| <code>mondata.nlm</code> | Allows you to monitor NetWare servers. |

Table 30-2 provides a brief description of the enterprise MIBs associated with the Management Agent for NetWare:

Table 30-2 List of Enterprise MIBs associated with the Management Agent for NetWare

| MIB Name | Description |
|---------------------------|---|
| <code>ndstrap.mib</code> | A Novell proprietary MIB designed to capture Novell eDirectory events and forward them to SNMP Management console as SNMP traps. There are more than 130 traps currently in the MIB and new ones are being added as they are identified. |
| <code>nwalarm.mib</code> | A Novell proprietary MIB that handles all the NetWare Core OS alerts and forwards them as SNMP traps. It currently supports more than 375 traps and new ones are being added as they are identified. |
| <code>nwhostx.mib</code> | A Novell extension to RFC1514 (the Host Resources MIB). It adds devices and components that are specific to NetWare that were not directly included in RFC1514. |
| <code>nwserver.mib</code> | A Novell proprietary MIB that is the basis for NetWare Core OS management. More than 300 objects are identified in this MIB. Access to the parameters that can be set from the console for both GET and SET is defined. The MIB has several groups and tables for users, file systems, volumes, queues, Open Data-Link Interface™ (ODI™), set parameters, and so forth. |
| <code>nwtmsync.mib</code> | A Novell proprietary MIB that allows for SNMP management of <code>timesync.nlm</code> . It provides access to the list of time sources as well as time clients. You may also access the clock structure through this MIB. |

| MIB Name | Description |
|-------------|---|
| nwtrend.mib | A Novell proprietary MIB that keeps track of objects that are most useful when tracked over a period of time. For example, CPU utilization and packets received have limited value as static numbers, but when monitored at regular intervals for a period of time, they tell a great deal about what is happening on a server. This MIB also lets you set user-definable thresholds for the managed objects and will send SNMP traps when a threshold is exceeded. |
| rfc1514.mib | The Internet Standard Host Resources MIB. It defines general categories about a host machine, including physical components of the system such as disks, memory, CPU, printers, adapter cards, and so forth. |

30.1.2 Management Agent for Windows Server Files

Table 30-3 lists files that can be manually configured with a text editor to modify default results of the Management Agent for Windows:

Table 30-3 List of Management Agent for Windows files

| Management Agent for Windows Server .INI Files | Description |
|--|---|
| n_nttren.ini | Specifies the initial values for the trends and thresholds supported by the Management Agent for Windows. |
| nttrap.ini | Specifies settings to troubleshoot your Windows server that runs the Management Agent for Windows and settings to enable you to send Windows events to the management system as SNMP traps. |
| nthost.ini | Specifies the SNMP settings supported by the Management Agent for Windows. |
| n_ntfmw.ini | Allows you to specify IPX addresses that will be ignored and will not receive SNMP traps. |

Table 30-4 provides a brief description of the enterprise MIBs associated with the Management Agent for Windows. In addition, the Management Agent for Windows converts all Windows system, security, and application events to SNMP traps.

Table 30-4 List of Enterprise MIBs associated with the Management Agent for Windows

| MIB Name | Description |
|--------------|--|
| ntserver.mib | Gives minimal Windows system information like Server Name, OS, major and minor versions, time zone, remote and local volumes count, etc. |
| rfc1514.mib | The Internet Standard Host Resources MIB. It defines general categories about a host machine, including physical components of the system such as disks, memory, CPU, printers, adapter cards, and so forth. |
| nttrap.mib | A generic MIB based on RFC1514. Windows events that are converted into traps are forwarded to the Novell ZENworks Server Management network management system. |

| MIB Name | Description |
|-------------|---|
| nttrend.mib | A Novell proprietary MIB that keeps track of objects that are most useful when tracked over a period of time. For example, CPU utilization and packets received have limited value as static numbers, but when monitored at regular intervals for a period of time, they tell a great deal about what is happening on a server. This MIB also lets you set user-definable thresholds for the managed objects and will send SNMP traps when a threshold is exceeded. |

30.2 Customizing the Management Agent for NetWare

The Management Agent for NetWare installation process creates the `nma2.ncf` file (Novell NetWare 3.x and 4.x servers) or the `nma5.ncf` file (NetWare 5.x servers) in the `sys:\zfs_agnt\nma` directory. When the NetWare server is started, this file automatically loads all the NLM files required for the Management Agent for NetWare in a default configuration state. There are, however, several LOAD parameters that you can configure for each of the NLM files used with the agent.

You can configure your server to use these options by editing the `nma2.ncf` or `nma5.ncf` file on your server. Also, if your server is already running, you can unload any of these NLM files and then load them at the NetWare server console using any of the configuration parameters. You can configure these parameters at the NetWare server console or by using the NetWare remote console utility, RCONSOLEJ.

The sections that follow describe each of the command line parameters that you can configure for the Management Agent for NetWare.

- ♦ [Section 30.2.1, “servinst.nlm Load Parameters,” on page 1118](#)
- ♦ [Section 30.2.2, “hostmib.nlm Load Parameters,” on page 1119](#)
- ♦ [Section 30.2.3, “ntrend.nlm Load Parameters,” on page 1120](#)

30.2.1 servinst.nlm Load Parameters

`Servinst.nlm` implements the `nwserver.mib` NetWare Server MIB. You can load `servinst.nlm` at the command line with any or all of the following parameters:

```
LOAD SERVINST D, U=n, V, B=n H
```

Table 30-5 List of `servinst.nlm` Load parameters

| Parameter | Description |
|-----------|---|
| D | DisableSets: If this parameter is present, <code>servinst.nlm</code> does not allow SNMP SET commands for objects in <code>nwserver.mib</code> . Default: SETS enabled (subject to SNMP security). |

| Parameter | Description |
|-------------|--|
| U= <i>n</i> | UpdateInterval= <i>n</i> : Sets the list update interval to <i>n</i> (<i>n</i> is a value in seconds). This determines how often certain internal lists kept by <code>servinst.nlm</code> (such as volumes and queues) are updated. Set this parameter higher to minimize the number of CPU cycles used by <code>servinst.nlm</code> , or lower to guarantee immediate reporting of server status changes that affect the lists. Default: 300 seconds. |
| V | Verbose: Displays informational messages. Default: Off. |
| B= <i>n</i> | BuildUserListHour= <i>n</i> : The local time each day on a 24-hour clock (0 to 23) at which the <code>servinst.nlm</code> software builds a list of users that have access to the server. Default: 2 (2:00 AM). |
| H | Help: Displays help on command line parameters. If you use the H parameter, <code>servinst.nlm</code> displays the help messages and then exits. It does not remain loaded even if other parameters are entered on the command line. Default: Off. |

30.2.2 hostmib.nlm Load Parameters

`Hostmib.nlm` implements both the standard Host Resources MIB (RFC 1514) and the Novell extensions to the Host Resources MIB (`nwhostx.mib`). You can load `hostmib.nlm` at the command line with any or all of the following parameters:

```
LOAD HOSTMIB.NLM D, U=n, V, H
```

Table 30-6 List of `hostmib.nlm` Load parameters

| Parameter | Description |
|-------------|--|
| D | DisableSets: If this parameter is present, <code>hostmib.nlm</code> does not allow SNMP SET commands for objects in RFC1514.MIB or <code>nwhostx.mib</code> . Default: SETS enabled (subject to SNMP security). |
| U= <i>n</i> | UpdateInterval= <i>n</i> : Sets the list update interval to <i>n</i> (<i>n</i> is a value in seconds). This determines how often certain internal lists kept by <code>hostmib.nlm</code> are updated. Set this parameter higher to minimize the number of CPU cycles used by <code>hostmib.nlm</code> , or lower to guarantee immediate reporting of server status changes that affect the lists. Default: 60 seconds. |
| V | Verbose: Displays informational messages. Default: Off. |
| H | Help: Displays help on command line parameters. If you use the H parameter, <code>hostmib.nlm</code> displays the help messages and then exits. It does not remain loaded even if other parameters are entered on the command line. Default: Off. |

30.2.3 ntrend.nlm Load Parameters

Ntrend.nlm implements the Threshold and Trend MIB (nwtrend.mib).

When first loaded, ntrend.nlm automatically sets trends and thresholds for each monitored attribute according to the server's configuration from values stored in the ntrend.ini file (located in the sys:\etc directory). You can edit this file as described in [“Setting Default Trends and Thresholds” on page 959](#).

Thereafter, as configuration changes occur over time, ntrend.nlm adjusts to changes in the number and type of physical network interfaces, queues, volumes, and disks. Default thresholds are set only for important parameters. You can later use SNMP SET commands to set thresholds for parameters such as files read and packets in.

A trend file is created for each monitored attribute instance, even if trending is disabled for that object. The file header contains all the information from nwtControlTableEntry, and the rest of the file stores the sample history (if any). After a trend file is created, it exists until explicitly deleted by the operator, even if the monitored object (a queue, for example) no longer exists. When a monitored object no longer exists, the associated nwtControlStatus is recorded as invalid.

You can load ntrend.nlm at the command line with any or all of the following parameters:

```
LOAD NTREND D=dir, R, V, H
```

Table 30-7 List of ntrend.nlm Load parameters

| Parameter | Description |
|---------------|---|
| D= <i>dir</i> | Directory= <i>dir</i> : Enables you to specify the volume and directory where ntrend.nlm stores the history data files. Example: To use vol1:\test as the directory for trending files, enter the following command: load ntrend D=vol1:\test Default: sys:\ntrend. |
| R | Reset: Causes ntrend.nlm to discard all the old trending history data and restart the sampling. |
| V | Verbose: Displays informational messages. Default: Off. |
| H | Help: Displays help on command line parameters. Default: Off. |

30.3 Customizing the Management Agent for Windows Server

You can manually edit the files listed in [Table 30-8](#) to modify the default Management Agent for Windows configuration on a managed Windows server:

Table 30-8 Management Agent for Windows Server .INI files

| Management Agent for Windows Server .INI Files | Description |
|--|--|
| <code>n_ntfmw.ini</code> | Allows you to specify IPX addresses that will be ignored and will not receive SNMP traps. |
| <code>nttrap.ini</code> | <p>Specifies settings to troubleshoot your managed Windows servers and set trap filters to specify which Windows events are sent to the management system as SNMP traps.</p> <p>See “Controlling Alarm Generation” on page 963 for detailed information on configuring trap filters and trap generation.</p> <p>The Server Management Agent has been enhanced with the following features: Section 30.3.2, “Collecting Events from Custom Event Log Types,” on page 1121 and Section 30.3.3, “Specifying Negative Filter Conditions in the Nttrap.ini File,” on page 1122.</p> |
| <code>n_nttren.ini</code> | <p>Specifies the initial values for the trends and thresholds supported by the Management Agent for Windows.</p> <p>See “Setting Default Trends and Thresholds” on page 959 for detailed information on modifying default trends and thresholds.</p> |

30.3.1 Configuring the Management Agent for Windows Server

By default, the Management Agent for Windows sends traps to SNMP Management console on IPX networks broadcasting the 0x026 Service Advertising Protocol (SAP) ID. You can edit the `n_ntfmw.ini` file to include the IPX addresses of SNMP Management console that you do not want to include as trap targets.

To add the IPX address of a SNMP Management console to omit as an automatic trap recipient:

- 1 Open the `n_ntfmw.ini` file in a text editor.
- 2 Add the IPX address for omitted SNMP Management console using the following syntax:
`xxxxxxxx.yyyyyyyyyyyy`
where `xxxxxxxx` is the net address and `yyyyyyyyyyyy` is the node address, such as `01014044.00001B4DDAFD`.
- 3 Save the file and restart the Management Agent for Windows.

30.3.2 Collecting Events from Custom Event Log Types

The Server Management Agent for Windows collects events from Windows Event Log and converts these events to traps. The traps are forwarded to the site server.

You can now use the Server Management Agent for Windows to specify custom event log types to collect traps from DNS server, directory service log, file replication service log, etc. using the `nttrap.ini` file.

- 1 Open the `nttrap.ini` file from the `installation_path\zfs_agnt\ntagent\ini` directory.

- 2 In the Monitor Settings section, enable the event for custom log types according to the instructions in the file.
- 3 Save the `nttrap.ini` file.

30.3.3 Specifying Negative Filter Conditions in the Nttrap.ini File

The Server Management Agent for Windows now enables you to specify conditions with negatives.

To specify negative filter conditions in the `nttrap.ini` file:

- 1 Open the `nttrap.ini` file from the `installation_path\zfs_agnt\ntagent\ini` directory.
- 2 In the Available Filters section, and specify the negative filter conditions. For example, Collect events except event with the ID as 500.
- 3 In the Actual Filters section, enable these filters.
- 4 Save the `nttrap.ini` file.

For more information about filter conditions and examples in the `nttrap.ini` file, see [TID 10098619](http://support.novell.com/cgi-bin/search/searchtid.cgi?10098619.htm) (<http://support.novell.com/cgi-bin/search/searchtid.cgi?10098619.htm>) in the Novell Support Knowledgebase.

30.4 Third-Party Agent Configuration

Third-party SNMP agents require the following tasks to be completed before traps are received:

- ♦ [Section 30.4.1, “Ensuring that Traps Are Received,” on page 1122](#)
- ♦ [Section 30.4.2, “Integrating Vendor-Specific SNMP Traps,” on page 1122](#)

30.4.1 Ensuring that Traps Are Received

When configuring the SNMP agent or SNMP Remote Network Monitoring (RMON) agent on a network device, configure the agent's trap destination list (trap-target list) to include the Novell ZENworks Server Management management server station IP address or server name. Refer to the agent's documentation for information on configuring this. Novell ConsoleOne displays alarms for all devices that forward alarms to the management server.

If your network device is using the Management Agent for NetWare, Management Agent for Windows Server, NetWare Traffic Analysis[®] Agent[™], or the Traffic Analysis Agent for Windows, the agent's trap destination list is automatically configured for you. For information on configuring the trap destination list, see [Section 23.2, “Setting Up Discovery,” on page 897](#) for configuration information and [Chapter 29, “Understanding Traffic Analysis,” on page 1029](#) for information on configuring the RMON agents.

30.4.2 Integrating Vendor-Specific SNMP Traps

Before the Alarm Management System can process the alarm, you must include vendor-specific MIBs for the third-party SNMP agents in the management server MIB pool. You can further integrate third-party SNMP agents by annotating the trap definitions in the vendor MIB.

The Alarm Management System interprets ASN.1 annotations to trap definitions in a MIB to set the severity level and device status assigned to an alarm. The MIBs included with Novell ZENworks Server Management already include the proper annotations. The annotations provide detail on severity levels and device status to the Alarm Management System.

See [Chapter 26, “Using the MIB Tools,” on page 995](#) for information on adding a MIB to the management server's MIB pool and annotating third-party MIBs.

30.5 Advanced Trending Agent

This section contains the following information:

- ◆ [Section 30.5.1, “What Is the Advanced Trending Agent?,” on page 1123](#)
- ◆ [Section 30.5.2, “Configuring the Trend Variables,” on page 1123](#)
- ◆ [Section 30.5.3, “Configuring the Advanced Trending Agent on Linux,” on page 1124](#)
- ◆ [Section 30.5.4, “Configuring the Advanced Trending Agent on All Platforms,” on page 1125](#)
- ◆ [Section 30.5.5, “Quick Reference Table,” on page 1126](#)
- ◆ [Section 30.5.6, “Refreshing Configuration Settings,” on page 1127](#)
- ◆ [Section 30.5.7, “Installing the Advanced Trending Agent,” on page 1128](#)

30.5.1 What Is the Advanced Trending Agent?

The Advanced Trending Agent is an application that gathers and stores the trend data (historic data) for any parameter instrumented by an SNMP agent, if it is defined by a MIB variable and not just pre-configured MIB variables.

The Advanced Trending Agent functionality is available on NetWare, Windows, and Linux platforms.

The Advanced Trending Agent periodically collects and stores the samples of the configured MIB variables. The collected data is exposed through an SNMP interface. This data can be used to view the long-term trend graphs of the parameters.

You can configure the Advanced Trending Agent using a configuration file or through SNMP interface. The Advanced Trending Agent includes a command line utility, which reads the configuration file for any updates. The utility also resets the data collected and starts collecting new data. You can also configure the Advanced Trending Agent with threshold values for the parameters and generate a trap when the value exceeds the threshold value.

30.5.2 Configuring the Trend Variables

The Advanced Trending Agent starts trending on MIB variables specified in a configuration file.

- ◆ **On NetWare:** `installation_path/advtrend/advtrend.ini`
- ◆ **On Windows:** `installation_path/advtrend/ini/advtrend.ini`
- ◆ **On Linux:** `/etc/opt/novell/Novell_ZENworks/zfs-mms-advtrend.conf`

You can modify the contents to add new parameters to trend upon or modify existing ones. The configuration file follows a standard INI format.

The section name can contain any value and is typically a meaningful name for the parameter being configured. All the section names in a given file must be distinct.

IMPORTANT: You cannot change the name of the section. If you do so, the data for the parameters is lost and new data is collected.

You cannot change the names of the keys. Each key represents a particular configuration that defines the trending activity for that parameter. There are certain mandatory keys that must be defined for any section. If these keys are not present, the Advanced Trending Agent ignores the entire section.

In a set of the configuration keys, some of the keys are treated as mandatory keys. That particular Configuration entry is read-only if the data for the Mandatory configuration keys is supplied or if the particular configuration entry is discarded.

A typical section looks like this:

```
[Interface In-Octets]
MIBVariable=1.3.6.1.2.1.2.2.1.10.1
MIBVariable Type=1
MIBVariableTrendingStatus=1
MIBVariableDisplayName=Interfaces
SampleInterval=5
SampleType=1
Buckets=80
ThresholdRising=100
ThresholdFalling=60
AlarmType=1
AlarmSeverity=4
ThresholdAlarmGenerationStatus=1
IdentifierOID={1.3.6.1.2.1.2.2.1.10.3}
Filters=2
F1={I1=MS TCP Loopback interface}
F2={I1=3COM 3C920 Integrated Fast Ethernet Controller 3C905C-TX
Compatible) - Packet Scheduler Miniport}
FilterType=1
```

You must specify the following mandatory keys:

MIBVariable (key 1)
SampleInterval (key 5)
ThresholdRising (key 8, only if ThresholdAlarmGenerationStatus=1)
ThresholdFalling (key 9, only if ThresholdAlarmGenerationStatus=1)
AlarmType1 (key 10, only if ThresholdAlarmGenerationStatus=1)

For more information, see [Section 30.5.5, “Quick Reference Table,” on page 1126](#).

30.5.3 Configuring the Advanced Trending Agent on Linux

Before starting the Advanced Trending Agent on Linux, ensure that you have completed the following tasks:

- ◆ [“Enabling the SNMPv1/SNMPv2c Access for the Local Host” on page 1125](#)
- ◆ [“Configuring the Trap Sink for Sending SNMP v1 Traps to the Site Server” on page 1125](#)

- ♦ “Changing the Default Community String Used by the Advanced Trending Agent” on page 1125

Enabling the SNMPv1/SNMPv2c Access for the Local Host

You must configure the `snmpd` master agent so that it allows SNMPv1/v2c read-only access to the Advanced Trending Agent.

- 1 Edit the `snmpd.conf` file that is used by the `snmpd` master agent.
- 2 Ensure that the access control settings allow SNMPv1/SNMPv2c access to the local host on the entire OID tree (starting with .1)

For more information, refer to the question, How do I configure access control in the [SNMP Frequently Asked Question \(http://www.net-snmp.org/FAQ.html\)](http://www.net-snmp.org/FAQ.html).

Configuring the Trap Sink for Sending SNMP v1 Traps to the Site Server

You must configure the `snmpd` master agent so that the traps generated by the Advanced Trending Agent are sent to the site server.

- 1 Edit the `snmpd.conf` file.
- 2 Add a line to configure a trap-sink pointing to the site server.

Refer to `man snmpd.conf (5)` man page for the specific syntax to configure a trap sink.

Changing the Default Community String Used by the Advanced Trending Agent

You must configure the Advanced Trending Agent to use the same community string you have used to allow access in the [Enabling the SNMPv1/SNMPv2c Access for the Local Host](#) section.

- 1 Edit the `/etc/opt/novell/Novell_ZENworks/zfs-mms-advtrend.conf` configuration file.
- 2 In the Agent Config section, set the value of the `ReadCommunity` key to the appropriate community string for which you have allowed SNMPv1/SNMPv2c access.

30.5.4 Configuring the Advanced Trending Agent on All Platforms

To configure the Advanced Trending Agent, you must change the default purge interval.

The Advanced Trending Agent automatically trends on certain available MIB variables or specific instances of some SNMP tables. If these instances become unavailable, you can clean up the resources associated with the trending. You must configure the purge interval the Advanced Trending Agent should wait before this occurs.

- 1 Edit the configuration file:
 - ♦ **On NetWare:** `installation_path/advtrend/advtrend.ini`
 - ♦ **On Windows:** `installation_path/advtrend/ini/advtrend.ini`
 - ♦ **On Linux:** `/etc/opt/novell/Novell_ZENworks/zfs-mms-advtrend.conf`
- 2 In the Agent Config section, set the value of `PurgeInterval` key to a value in seconds.

30.5.5 Quick Reference Table

Table 30-9 *Values of the keys*

| Key Name | Explanation | Allowed Values | Mandatory | Default Value |
|--------------------------------|---|-----------------------------|---|---------------|
| MIBVariable | The variable to trend on | Any integer valued SNMP OID | Yes | None |
| MIBVariableType | Whether the variable is scalar or whether it is a column in an SNMP table | 1 (Scalar), 2 (Columnar) | No | 1 |
| MIBVariableTrendingStatus | Whether the trending is enabled on this or disabled | 1 (Enabled), 2 (Disabled) | No | 2 |
| MIBVariableDisplayName | The name to refer to this variable from console Views | Any string | No | Object |
| SampleInterval | The time interval in seconds indicating when to take the sample | Positive integer | Yes | None |
| SampleType | Whether samples stored must be absolute values or deltas | 1 (Absolute), 2 (Delta) | - | - |
| Buckets | Number of samples to store for a variable | Positive integer | No | 50 |
| ThresholdAlarmGenerationStatus | Whether to send traps or not | 1 (Enabled), 2 (Disabled) | No | 2 |
| ThresholdRising | If the sample value exceeds the first time and the AlarmType value is set to rising, send a trap | Integer | Yes, if ThresholdAlarmGenerationStatus is enabled | None |
| ThresholdFalling | If the sample value falls exceeds the first time and the AlarmType value is set to Falling, send a trap | Integer | Yes, if ThresholdAlarmGenerationStatus is enabled | None |

| Key Name | Explanation | Allowed Values | Mandatory | Default Value |
|---------------|--|--|---|---------------|
| AlarmType | Whether the trap should be generated for exceeding the ThresholdRising or ThresholdFalling value | 1 (Rising), 2 (Falling) | Yes, if ThresholdAlarmGenerationStatus is enabled | None |
| AlarmSeverity | Criticality of the trap that is sent. | 1 Severe 2 Major 3 Minor 4 Information 5 Unknown | No | 3 |
| IdentifierOID | For an OID of MIBVariableType 2, this uniquely identifies the rows in a table. | A formatted string | No | None |
| Filters | Number of filter conditions specified. | Integer between 1 and 5 | No | None |
| FilterType | Whether a SNMP row that matches a filter must be trended or ignored. | 1 Inclusive (include), 2 Exclusive (ignore) | No | 1 |

30.5.6 Refreshing Configuration Settings

The Advanced Trending Agent enables you to dynamically change the parameters of the trending activity for one or more MIB variables. You can add new entries or remove existing entries that are being trended in the configuration file.

After you make the changes, you must enter the `advtrend` command for the Advanced Trending Agent to retrieve the latest information. The `advtrend` command is available in the following locations:

- ◆ **On Windows:** `install_directory/zfs_agnt`
- ◆ **On Linux:** `/opt/novell/novell zenworks/bin`

Use the `advtrend` command to perform the following operations:

- ◆ [“Reset the Information” on page 1128](#)
- ◆ [“Adding a New Parameter for Trending” on page 1128](#)
- ◆ [“Removing an Existing Trended Parameter” on page 1128](#)
- ◆ [“Modifying the Values of the Keys for an Existing Parameter” on page 1128](#)

Reset the Information

If you want to restart trending the parameters:

- 1 Make changes to the configuration file to add, modify, or delete any configuration sections.
- 2 At the server prompt, enter `advtrend reset`.

The new entries will replace existing entries according to the configuration settings. The trending will start from the beginning.

WARNING: Once you reset the information, the old data will be lost.

Adding a New Parameter for Trending

- 1 Edit the configuration file to include a new section for your parameter.
- 2 Define the required keys in the new section.
- 3 At the server prompt on NetWare enter `advtrend reread`, and for Windows or Linux, enter `advtrend read_cfg`.

The new entries to be trended are appended.

Removing an Existing Trended Parameter

1. Edit the configuration file to remove the section representing the parameter being trended.
2. At the server prompt on NetWare enter `advtrend reread`, and for Windows or Linux, enter `advtrend read_cfg`.

The entries to be trended for this configuration section are removed.

Modifying the Values of the Keys for an Existing Parameter

If you want to modify the values for the keys of an MIB variable being trended, do the following:

- 1 Edit the configuration file to change the values of the keys.
- 2 At the server prompt, on NetWare enter `advtrend reread`, and for Windows or Linux, enter `advtrend read_cfg`.

The Advanced Trending Agent will now use the new parameters.

If you modify the following parameters, the previous data will be removed:

```
MIBVariable
MIBVariableType
SampleInterval
SampleType
IdentifierOID
```

30.5.7 Installing the Advanced Trending Agent

For more information on how to install the Advanced Trending Agent, see the *Novell ZENworks 7 Server Management Installation Guide*.

30.6 Management and Monitoring Services for Linux

The Management and Monitoring Services component of Novell ZENworks Server Management for Linux now provides you with the ability to centrally manage and administer the Linux servers on your network.

You can view all the real time statistical information and the historical information as views from Novell ConsoleOne®. You can manage traps generated for important events, obtain historical information about the Linux servers and view all the information as views from Novell ConsoleOne and generate reports about the overall health of your Linux server.

The Linux Management Agent of Management and Monitoring Services allows you to manage and monitor all the information about your Linux servers.

- ◆ [Section 30.6.1, “Providing Real Time Statistical Information,” on page 1129](#)
- ◆ [Section 30.6.2, “Generating Traps for System Events,” on page 1130](#)
- ◆ [Section 30.6.3, “Providing History Collection Information,” on page 1132](#)
- ◆ [Section 30.6.4, “Linux Management Views,” on page 1132](#)
- ◆ [Section 30.6.5, “Linux Server Health Reports,” on page 1133](#)

30.6.1 Providing Real Time Statistical Information

The Linux Management Agent is an SNMP agent that gathers and provides real time statistical information on some of the critical resources of the server.

The Linux Management Agent provides the following statistical information:

- ◆ Processor statistics
- ◆ System level memory usage
- ◆ Disk usage and statistics
- ◆ Partition usage statistics
- ◆ Init.d services statistics
- ◆ Process information
- ◆ User login statistics
- ◆ Kernel cache information
- ◆ Disk partition statistics
- ◆ System paging and swapping
- ◆ Interrupt Statistics on various processors

Customizing the Configuration File

Using the `zfs-mms-servinst.conf` file, you can customize the Linux Management Agent to monitor the real-time statistics. You can change the list of services that are automatically monitored in the `zfs-mms-servinst.conf` file.

- 1 Open the `zfs-mms-servinst.conf` file in the `/etc/opt/novell/novell zenworks` directory.

- 2 In the Monitored Services section, add or remove the service name from list of services.
- 3 Save the `zfs-mms-servinst.conf` file.
- 4 Restart the `snmpd`.

30.6.2 Generating Traps for System Events

The Linux Management Agent gathers important and critical events about the running server processes. Based on the criticality of the events, traps are generated and forwarded to the site server. The Linux Management Agent includes the `novell-trapd` service that monitor various categories of system logs for these events. Some of the categories of events gathered from the system logs include:

- ◆ User login
- ◆ Root login failures
- ◆ Service start/stop/restart
- ◆ Important kernel events logged

Customizing the Configuration File

Using the `zfs-mms-log2trap.conf` file, you can customize the Linux Management Agent to monitor the events for which the traps are generated.

You can do the following methods:

Changing the list of hosts where the traps are forwarded

- 1 Open the `zfs-mms-log2trap.conf` file in the `/etc/opt/novell/novell_zenworks` directory.
- 2 In the `TrapTargets` section, add the IP addresses of the hosts where you would like the traps to be sent. In the same line, add the SNMP community string that is used to send the traps.
- 3 Save the `zfs-mms-log2trap.conf` file.
- 4 Restart the `novell-trapd` service.

Preventing the trap generation for certain system log messages

Using the `zfs-mms-log2trap.conf` file you can prevent the traps from being generated for certain system log messages. You must create sections within the configuration file and include the keys in [Table 30-10](#):

Table 30-10 List of keys and its possible values to be used in the system log messages

| Key Name | Description | Possible Values |
|----------------------|---|---|
| Severity | The severity of the log message to be ignored. | <ul style="list-style-type: none"> ◆ CRITICAL ◆ MAJOR ◆ MINOR ◆ INFORMATIONAL ◆ MISCELLANEOUS ◆ ALL |
| Constraints | Defines how a log message should be searched for a substring. If the substring satisfies the conditions specified, this log message is ignored. | <p>The possible values are an optional CASE_COMPARE and one of CONTAINS or ENDS_WITH or BEGINS_WITH.</p> <p>The optional CASE_COMPARE specifies if a case-sensitive search is required on the string. The other values specify whether the string must be contained in the log, or the end of the log or the beginning of the log respectively.</p> |
| string <i>number</i> | Defines a string that is searched should be searched. You can define any number of such keys, by incrementing the value of <i>number</i> . A message that matches any one of these values is ignored. | Any string. |

Example: Rejecting all log messages containing the string telnet

- 1 Open the `zfs-mms-log2trap.conf` file in the `/etc/opt/novell/Novell ZENworks` directory.
- 2 Create a section in the configuration file. For example `[Telnet Reject]`.
- 3 Add the following keys and values:


```
severity = ALL
constraints = CASE_COMPARE
string1 = telnet
```
- 4 Search for the section `_Reject Setting Names_`, and add the section name. For example, `Telnet Reject`.
- 5 Save the configuration file.
- 6 Restart the `novell-trapd` service.

Example: Rejecting all messages with severity value as Info containing the string 127.0.0.1 or localhost

- 1 Open the `zfs-mms-log2trap.conf` file in the `/etc/opt/novell/Novell ZENworks` directory.
- 2 Create a section in the configuration file. For example `[Localhost]`.
- 3 Add the following keys and values:

```
severity = INFO
constraints = CASE_COMPARE
string1 = 127.0.0.1
string2 = localhost
```

- 4 Search for the section `_Reject Setting Names_`, and add the section name. For example, `Localhost`.
- 5 Save the configuration file.
- 6 Restart the `novell-trapd` service.

30.6.3 Providing History Collection Information

The Linux Management Agent implements the SNMP instrumentation for the Linux Operating System for additional information. Using the SNMP interface information the following attributes of the Linux Operating System can be obtained:

- ◆ CPU Utilization
- ◆ Used Memory and Used Swap Size
- ◆ Disk Free Space
- ◆ Network Interface Statistics (incoming and outgoing packet count)
- ◆ Disk Reads and Writes
- ◆ Logged in User Count

Using Advanced Trending agent you can configure and control periodic statistical sampling for any of the above parameters. Using the Advanced Trending Agent you can also set thresholds to generate traps for the above attributes.

For more information on the Advanced Trending Agent, see [Section 30.5, “Advanced Trending Agent,” on page 1123](#).

30.6.4 Linux Management Views

Using the Unified View for Devices, you can list all the Linux devices on your network.

The following views are created in Linux to enable you to access more information about a Linux server:

- ◆ Node Summary
- ◆ Processor Statistics and Trend
- ◆ Storage Devices Summary and Trend
- ◆ Running Software
- ◆ Interrupt Summary Node
- ◆ Memory Statistics and Trend
- ◆ Network Interface Statistics and Trend
- ◆ Currently Logged-in Users and Trend
- ◆ Services Started by the Node

The Linux Management views are similar to the views on NetWare and Windows servers.

30.6.5 Linux Server Health Reports

A new profile called the Linux Server Profile is added to the list of profiles to enable you to generate reports. This reports enables you to obtain information about the overall health of your Linux server. You can use the Linux Server Health Report to generate reports daily, weekly, monthly.

Protocol Decodes Suites Supported by Novell ZENworks Server Management

Novell ZENworks[®] Server Management provides packet capture and decoding tools that help you analyze the network activity and identify the source of network problems. Capturing and decoding packets can help you troubleshoot network problems by giving you detailed information about segment activity. For details, see [“Capturing Packets” on page 1063](#) and [“Displaying Captured Packets” on page 1066](#).

This section provides information about decoding support provided by Novell ZENworks Server Management for the following protocol suites:

- ◆ [Section 31.1, “Novell NetWare Protocol Suite,” on page 1135](#)
- ◆ [Section 31.2, “Network File System Protocol Suite,” on page 1137](#)
- ◆ [Section 31.3, “Systems Network Architecture Protocol Suite,” on page 1137](#)
- ◆ [Section 31.4, “AppleTalk Protocol Suite,” on page 1138](#)
- ◆ [Section 31.5, “TCP/IP Protocol Suite,” on page 1139](#)

31.1 Novell NetWare Protocol Suite

Novell NetWare[®] contains a group of protocols that perform various functions in a Novell NetWare network. Each protocol in the Novell NetWare protocol suite works with the IPX[™] protocol. Novell ZENworks Server Management supports the following protocols in the Novell NetWare suite of protocols:

Table 31-1 Novell NetWare suite of protocols

| Novell NetWare Protocol | Description |
|-------------------------|---|
| BCAST | Novell NetWare Broadcast Message Notification. The protocol a Novell NetWare server uses to inform an idle workstation that a message is pending. This message appears on the top or bottom line of the monitor on DOS stations. |
| DIAG | Diagnostic Responder. A protocol used for connectivity testing and information gathering. By default, Novell NetWare clients use the Diagnostic Responder to reply to diagnostic requests. |
| IPX | Internetwork Packet Exchange [™] . A protocol that routes outgoing data packets across a network. Every Novell NetWare network has a unique address assigned when its servers are configured. IPX routers use this address to route packets through an internetwork. IPX makes routing decisions based on information compiled by the Routing Information Protocol (RIP). |

| Novell NetWare Protocol | Description |
|--------------------------------|---|
| LSP | Novell NetWare Lite™ Sideband Protocol. A connectionless (datagram) oriented protocol that operates as a sideband for Novell NetWare Lite Transport Protocol (NLTP) connections. |
| NBIOS | NetBIOS. An emulator that allows workstations to run applications that support IBM* NetBIOS calls. NetBIOS is the IBM standard protocol for applications developed to run peer-to-peer communications on token ring networks. |
| NCP™ | Novell NetWare Core Protocol™. A set of procedures that a file server operating system follows to accept and respond to workstation requests. An NCP exist for every service a workstation might request from a file server. Common requests handled by the NCP protocols include creating or deleting a file, manipulating directories and files, performing a directory listing, altering the bindery (drive mappings and security), and printing. |
| NDS® | The NDS protocol, called the Novell Directory Access Protocol (NDAP), is a wire protocol that allows Novell eDirectory to service client requests and to send client requests to other Novell eDirectory servers. NDAP is built based on NCP. |
| NLP | Novell NetWare Lite Protocol. A protocol that is an integral part of Novell NetWare Lite, which operates on top of the Novell IPX protocol. NLP is an application-layer and service-layer protocol that performs file system and print functions. NLP also uses NLTP, which is similar in function to the transport protocol used in NCP. |
| NLSP™ | Novell NetWare Link Services Protocol™. A link-state routing protocol designed for IPX internetworks. |
| RIP | Routing Information Protocol. A protocol that automates the process of updating routing tables. Routing is the process of moving network packets between separate networks. With RIP, when one router learns about changes in its routes, it broadcasts this information to neighboring routers so they can update their routing tables. As a result, if a network component fails (such as a router or a phone line), the other network components can inform each other of alternate routes. When the faulty component is repaired, the network changes back to the previous condition. |
| SAP | Service Advertising Protocol. A protocol that lets Novell NetWare servers advertise their services by name and type. A workstation can broadcast a request to find all services available or a specific service closest to the client. |
| SER | Novell Serialization (Copy Protection) Packets. Packets that Novell NetWare servers send to other Novell NetWare servers to ensure that each server has a unique serial number. |
| SNMP | Simple Network Management Protocol. An application-layer protocol designed to facilitate the exchange of management information between network devices. By using SNMP to access management information data (such as packets per second and network error rates), network administrators can easily manage network performance and find and solve network problems. |

| Novell NetWare Protocol | Description |
|-------------------------|--|
| SPX™ | <p>Sequenced Packet Exchange™. A connection-oriented transport protocol that monitors network transmissions to ensure successful delivery of packets. SPX enhances the IPX protocol by supervising data sent across the network. SPX can track data transmissions consisting of a series of separate packets.</p> <p>SPX also requests acknowledgments from and returns acknowledgments to a communications partner, ensuring successful data delivery. If an acknowledgment request brings no response within a specified time, SPX retransmits the request. After a reasonable number of retransmissions fail to return a positive acknowledgment, SPX assumes the connection has failed and reports the error.</p> <p>The Novell NetWare print server uses SPX.</p> |
| WDOG | <p>Watchdog. A maintenance protocol provided with Novell NetWare. Watchdog monitors stations that are logged in to a Novell NetWare server. Watchdog determines whether the Novell NetWare shells are still operating and, if not, releases the connection.</p> |

31.2 Network File System Protocol Suite

The Network File System (NFS) suite of protocols is described in [Table 31-2](#):

Table 31-2 *Network File System (NFS) suite of protocols*

| Network File System Protocol | Description |
|------------------------------|--|
| MOUNT | <p>The MOUNT protocol, used in conjunction with NFS, performs operating system-specific functions that allow NFS clients to attach remote directory trees to a point within the local file system.</p> |
| NFS | <p>Network File System. This protocol provides transparent remote access to shared file systems across networks. NFS uses Remote Procedure Call (RPC) and is machine, operating system, network architecture, and transport protocol independent.</p> |
| PORTMAP | <p>The PORTMAP protocol converts RPC program numbers into Transmission Control Protocol/User Datagram Protocol (TCP/UDP) port numbers. When a client wants to make an RPC call to a given program number, it will first contact PORTMAP on the remote machine to determine the port number where RPC packets should be sent.</p> |
| RPC | <p>Remote Procedure Call. This protocol allows a program on one computer to make a subroutine call on a remote computer. Every subroutine or remote procedure is identified by a unique program number.</p> |

31.3 Systems Network Architecture Protocol Suite

The Systems Network Architecture (SNA) suite of protocols is described in [Table 31-3](#):

Table 31-3 *Systems Network Architecture (SNA) suite of protocols*

| Systems Network Architecture Protocol | Description |
|--|--|
| RH | Request/Response Header. This protocol carries the SNA Request/Response Units as its payload. |
| RU | Request/Response Unit. An SNA client uses this protocol to communicate with an SNA server. |
| TH | Transmission Header. This protocol runs on a data link layer and serves as the transmission layer for an SNA Path Information Unit. |
| XID | Exchange Station Identification. An SNA node uses this protocol to check whether its peer SNA node is ready for communication and to exchange its station details with it. |

31.4 AppleTalk Protocol Suite

The AppleTalk* and AppleTalk-related suite of protocols is described in [Table 31-4](#):

Table 31-4 *AppleTalk and AppleTalk-related suite of protocols*

| AppleTalk Protocol | Description |
|---------------------------|--|
| AARP | AppleTalk Address Resolution Protocol. An AppleTalk protocol that reconciles addressing differences between a data link protocol and the rest of a protocol family. For example, by resolving the differences between an Ethernet addressing scheme and the AppleTalk addressing scheme, AARP facilitates the transport of datagram delivery protocol (DDP) packets over a high-speed EtherTalk* connection. |
| ADSP | AppleTalk Data Stream Protocol. A connection-oriented protocol that provides a reliable, full-duplex, byte stream service between any two sockets in an AppleTalk internetwork. ADSP ensures sequential, duplicate-free delivery of data over its connections. |
| AEP | AppleTalk Echo Protocol. A simple protocol that allows a node to send a packet to any other node in an AppleTalk internetwork and receive an echoed copy of that packet in return. |
| AFP | AppleTalk Filing Protocol. A presentation layer protocol that allows users to share data files and applications that reside in an AppleTalk shared resource, such as a file server. |
| ASP | AppleTalk Session Protocol. A general, all-purpose protocol that uses the services of the AppleTalk Transaction Protocol (ATP) to provide session establishment, maintenance, and tear-down, along with request sequencing. |
| ATP | AppleTalk Transaction Protocol. A transport protocol that provides a loss-free transaction service between sockets. This service allows exchanges between two socket clients in which one client requests the other to perform a particular task and report the results. ATP binds the request and response together to ensure the reliable exchange of request-response pairs. |

| AppleTalk Protocol | Description |
|--------------------|---|
| E-DDP | Extended Datagram Delivery Protocol. A datagram delivery protocol that uses an extended header. An extended header is required for packets that are transmitted from one network to another network within an AppleTalk Internet. |
| ELAP | EtherTalk Link Access Protocol. The link-access protocol used in an EtherTalk network. It is built on the top of the standard Ethernet data link layer. |
| NBP | Name Binding Protocol. A transport layer protocol that translates a character string name into the internetwork address of the corresponding socket client. NBP enables AppleTalk protocols to understand user-defined zones and device names by providing and maintaining translation tables that map these names to corresponding socket addresses. |
| PAP | Printer Access Protocol. This protocol manages interaction between workstations and print servers. It handles connection setup, maintenance, and termination. It can also handle data transfer. |
| RTMP | Routing Table Maintenance Protocol. This AppleTalk protocol establishes and maintains the routing information that is required by internetwork routers to route datagrams from any source socket to any destination socket in the internetwork. Using RTMP, internetwork routers dynamically maintain routing tables to reflect changes in internetwork topology. |
| S-DDP | Short Datagram Delivery Protocol. A DDP that uses a short header. A short header is often used for packets whose source and destination sockets are within the boundaries of a single AppleTalk network. |
| ZIP | Zone Information Protocol. A protocol that maintains up-to-date routing information across the internetwork. |

31.5 TCP/IP Protocol Suite

The TCP/IP suite of protocols is described in [Table 31-5](#):

Table 31-5 *TCP/IP suite of protocols*

| TCP/IP Protocol | Description |
|-----------------|---|
| ARP | <p>Address Resolution Protocol. A protocol used by a host to determine the hardware address of another host. A TCP/IP system contains a table that maps IP addresses to the hardware addresses of the different hosts and routers on the internetwork. This table works in much the same way as a host table, translating an IP address to an Ethernet address. Unlike the host table, however, the ARP table is not usually maintained by you or your network administrator. The ARP protocol creates entries in this table as needed.</p> <p>If the hardware address of the destination is not found in your station's ARP table, a broadcast is sent to every host on the network requesting the address. If that host is up and supports the ARP protocol, it receives the broadcast from your station and responds by sending its hardware address back to your station. This address is then added to your station's ARP table.</p> |
| IMAP | IMAP stands for Internet Message Access Protocol. It is a method of accessing electronic mail or bulletin board messages that are placed on a (possibly shared) mail server. It permits a "client" e-mail program to access remote message stores as if they were local. |

| TCP/IP Protocol | Description |
|------------------------|---|
| BOOTP | BootStrap Protocol. This protocol allows a diskless workstation to determine its IP address and other information without using the Reverse Address Resolution Protocol (RARP). |
| DHCP | Dynamic Host Configuration Protocol. This protocol supplies hosts with configuration parameters, leases dynamically allocated IP addresses, and acts as an enhancement to BOOTP. |
| DNS | Domain Name System. The distributed naming service used on the Internet. DNS provides a computer's IP address if domain names exist for the computer. |
| FTP | File Transfer Protocol. TCP/IP application-layer protocol that supports file transfers. |
| HTTP | Hypertext Transfer Protocol. An application-layer protocol that Web browsers and Web servers use to communicate with each other. |
| ICMP | Internet Control Message Protocol. A protocol that works with IP to provide routing efficiency and error information. ICMP is part of the TCP/IP protocol suite. Because IP is connectionless, it cannot detect anomalous internetwork conditions. ICMP works with IP to provide TCP or other upper-layer protocols with this information. |
| IGMP | Internet Group Management Protocol. A protocol used by IP hosts to report their multicast group memberships to routers. The protocol is also used to query routers on memberships and to generate reports on group membership. Termination of group membership can be quickly reported using this protocol. |
| IP | <p>Internet Protocol. A protocol that provides connectionless, nonguaranteed delivery of transport layer packets (also called transport protocol data units or TPDU) across an internetwork. IP is part of the TCP/IP protocol suite.</p> <p>IP can fragment TPDU) into smaller parts, if necessary, and then reassemble them at an intermediate station (usually a router) or at their destination host.</p> <p>Each TPDU or fragment is fitted with an IP header and transmitted as a packet by lower-layer protocols. IP moves datagrams through the internetwork, one hop at a time. If a TPDU fragment arrives at its destination out of order, IP reassembles the fragments, in sequence, at the destination.</p> |
| LDAP | Lightweight Directory Access Protocol. This protocol provides access to the x.500 Directory while not incurring the resource requirements of the Directory Access Protocol (DAP). LDAP is specifically targeted at simple management applications and browser applications that provide read/write interactive access to the x.500 Directory, and is intended to be a complement to the DAP itself. |
| NFS | The Network File System (NFS) protocol provides transparent remote access to shared files across networks. The NFS protocol is designed to be portable across different machines, operating systems, network architectures, and transport protocols. This portability is achieved through the use of Remote Procedure Call (RPC) primitives built on top of an eXternal Data Representation (XDR). |
| NTP | Network Time Protocol. A protocol used to synchronize timekeeping among a set of distributed time servers and clients. It is used to convey timekeeping information in a hierarchical method from servers to clients. It is also used to cross-check clocks and control errors due to equipment or propagation failures. |

| TCP/IP Protocol | Description |
|-----------------|---|
| NWIP | <p>Novell NetWare/IP. Allows total or partial replacement of the IPX transport subsystem with the industry-standard TCP/IP subsystem, in a Novell NetWare network. The following constitute the core components of the technology:</p> <ul style="list-style-type: none"> ◆ Communication between the Novell NetWare/IP server and the Domain SAP/RIP Service (DSS) for <ul style="list-style-type: none"> - Retrieval of configuration parameters - Registration of SAP and RIP information - SAP/RIP database synchronization ◆ Synchronization of the Novell NetWare/IP server with the DSS database with respect to SAP/RIP information ◆ Communication between secondary DSS and primary DSS to synchronize the SAP/RIP database on the two servers |
| OSPF | <p>Open Shortest Path First. A protocol in the TCP/IP protocol suite is an interior gateway protocol algorithm and is proposed as a standard for the Internet. OSPF incorporates least-cost routing, multipath routing, load balancing, and efficient bandwidth utilization.</p> |
| POP3 | <p>Post Office Protocol 3. A protocol used for interacting with a central mailbox server. It is a client/server protocol used to receive e-mail. The protocol holds the e-mail messages in the Internet server. Periodically, you can download the messages from the server.</p> |
| RARP | <p>Reverse Address Resolution Protocol. A protocol in the TCP/IP protocol suite that is used to determine a software address based on a hardware address. This protocol is often used by diskless workstations during startup.</p> |
| RIP | <p>Routing Information Protocol. A protocol in the Novell NetWare protocol suite that automates the process of updating routing tables. Routing is the process of moving network packets between separate networks. With RIP, when one router learns about changes in its routes, it broadcasts this information to neighboring routers so they can update their routing tables. As a result of RIP, if a network component fails (such as a router or a phone line), the other network components can inform each other of alternate routes. When the faulty component is repaired, the network changes back to the previous condition.</p> |
| SSL | <p>SSL is an open, nonproprietary protocol. It has been submitted to the W3 Consortium (W3C) working group on security for consideration as a standard security approach for World Wide Web browsers and servers on the Internet.</p> |
| SLP | <p>Service Location Protocol. This protocol provides a scalable framework for the discovery and selection of network services. Using this protocol, computers using the Internet no longer need as many static configurations of network services for network-based applications.</p> |
| SMTP | <p>Simple Mail Transfer Protocol. The application layer protocol that e-mail clients and servers use to exchange e-mail messages with each other.</p> |

| TCP/IP Protocol | Description |
|-----------------|---|
| SNMP | <p data-bbox="558 260 1419 401">Simple Network Management Protocol. A protocol in the TCP/IP protocol suite that enables you to monitor a network from a single network management station called an SNMP Manager. From an SNMP Manager, you can make inquiries to another network device called the SNMP Agent. The SNMP Agent can be a TCP/IP host, router, terminal server, or another SNMP Manager.</p> <p data-bbox="558 428 1419 625">The information you can request from an SNMP Agent is contained in the MIB of that TCP/IP host. RFC 1066 (http://www.isi.edu/in-notes/rfc1066.txt) (Internet standard MIB) defines the types of objects that can be in an SNMP Agent MIB. These objects include network and hardware addresses, counters, and statistics, as well as routing and Address Resolution Protocol tables. Different vendors might not support all data types within their MIB or might include other information not defined within the RFC.</p> |
| TCP | <p data-bbox="558 653 1386 737">Transmission Control Protocol. This primary Internet transport protocol accepts messages of any length from an upper-layer protocol and provides full-duplex, acknowledged, connection-oriented, flow-controlled transport.</p> |
| TELNET | <p data-bbox="558 764 1354 785">Protocol in the TCP/IP suite that governs character-oriented terminal traffic.</p> |
| TFTP | <p data-bbox="558 812 1338 863">Trivial File Transfer Protocol. TCP/IP protocol commonly used for software downloads.</p> |
| UDP | <p data-bbox="558 890 1419 1087">User Datagram Protocol. A protocol similar to TCP that provides connectionless, nonguaranteed transport services. UDP accepts and transports datagrams from an upper-layer protocol. Unburdened by the overhead of establishing and removing connections, controlling data flow, and performing other TCP functions, UDP usually provides a faster data conduit than TCP. For these reasons, and because it is easier to implement, UDP is the transport method of choice for many upper-layer protocols.</p> |

Novell ZENworks Management and Monitoring Services Database

32

Novell ZENworks[®] Server Management provides a centralized Common Information Model (CIM)-compliant Sybase* database on the Management and Monitoring Services management server. The database serves as a repository for server and network data that can be displayed or formatted in various ways to provide you with exactly the information you need to manage your network.

The following sections provide information on understanding and using the Novell ZENworks database:

- ♦ [Section 32.1, “Understanding the Novell ZENworks Server Management Database,” on page 1143](#)
- ♦ [Section 32.2, “Backing Up the Topology/Alarm Database,” on page 1144](#)
- ♦ [Section 32.3, “Changing Database Passwords,” on page 1144](#)
- ♦ [Section 32.4, “Emptying the Database,” on page 1144](#)

32.1 Understanding the Novell ZENworks Server Management Database

The Novell ZENworks Server Management database consists of files located in the `\novell zenworks\mms\db` directory on the management server. The Novell ZENworks Server Management data is stored in the following logical database:

- ♦ Topology/alarm database containing topology, alarms, and map information associated with the following files:
 - ♦ `Mw.db`
 - ♦ `Mw1.db`
 - ♦ `Mw2.db`
 - ♦ `Mw3.db`

The `mw.log` file in the `\novell zenworks\mms\db` subdirectory saves your transaction information with the database files.

- ♦ [Section 32.1.1, “Running the Database,” on page 1143](#)
- ♦ [Section 32.1.2, “Database Caching,” on page 1144](#)

32.1.1 Running the Database

The database is run using the `mgmt dbs.ncf` file (located in the `system` directory on a server volume), which is executed from `autoexec.ncf`.

IMPORTANT: Ensure that the database is running as long as the Novell ZENworks Server Management services are running.

32.1.2 Database Caching

Increasing the database cache improves the database performance. The default database cache size is 48 MB. You can increase the cache size to an optimum level depending on the server memory. To increase the cache size, modify the `-c` option in `sys:\system\mgmt\dfs.ncf`. For example, `-c 64M` sets the cache size to 64 MB. Reload the database after modifying the cache size.

32.2 Backing Up the Topology/Alarm Database

You should plan to regularly back up the Novell ZENworks Server Management database. In Novell ConsoleOne, follow this procedure to back up the topology/alarm database:

- 1 Right-click the Site Server object > select *Properties*.
- 2 Select the *Database Administration* tab.
- 3 Enter the path of the directory to back up.

You can back up the database files to any volume on the management server only.

- 4 Click *Apply*.

Novell ZENworks Server Management sends a remote SQL command to store the file. The four `MW*.DB` and `mw.log` files are copied to the backup directory.

32.3 Changing Database Passwords

Novell ZENworks Server Management allows you to access the topology/alarm database at three different levels: Administrator account, Updater account, and Reader account. You can set passwords for any of the three different user accounts.

In Novell ConsoleOne, follow this procedure to modify the database passwords:

- 1 Right-click the Site Server object > select *Properties*.
- 2 Select the *Change Database Passwords* tab.
- 3 Enter the new passwords and confirm.
- 4 Click *Apply*.

Novell ZENworks Server Management sends a remote SQL command to change the passwords of appropriate user objects in the database. The passwords are also stored in the Novell eDirectory

32.4 Emptying the Database

Novell ZENworks Server Management enables you to empty the current database you are using and use a fresh database. The database files that you are using currently are located in the `novell zenworks\mms\db` directory. To empty the database:

- 1 Select all the files from the `novell zenworks\mms\emptydb` directory.
- 2 Copy all the files from the `novell zenworks\mms\emptydb` directory to the `novell zenworks\mms\db` directory.

Using Reports in Management and Monitoring Services

33

The Novell ZENworks® Server Management Management and Monitoring Services provide the following predefined reports:

- ♦ Topology Reports
- ♦ Alarm Reports
- ♦ Health Reports

The following sections describe the available reports and provide procedures for customizing and generating the reports:

- ♦ [Section 33.1, “Understanding Management and Monitoring Services Reports,” on page 1145](#)
- ♦ [Section 33.2, “Managing Reporting,” on page 1151](#)

33.1 Understanding Management and Monitoring Services Reports

The following sections describe each predefined report available in Management and Monitoring Services:

- ♦ [Section 33.1.1, “About the Topology Reports,” on page 1145](#)
- ♦ [Section 33.1.2, “About the Alarm Reports,” on page 1148](#)
- ♦ [Section 33.1.3, “About the Health Reports,” on page 1149](#)

33.1.1 About the Topology Reports

The topology reports provide information about the topology of a selected Novell ZENworks Server Management Site, segment or custom atlas container. The site-level reports provide details about the discovered devices on each segment in the Novell ZENworks Server Management site. The segment-level reports provide information about the discovered devices on the selected network segment. The custom atlas container report provides information about the computer system in the select container.

Prior to generating the reports, you will need to perform a few operations. For more information see [“Prerequisites for Generating the Reports” on page 1146](#).

There are five predefined topology reports:

- ♦ [“Computer Systems by Segment Report” on page 1146](#)
- ♦ [“NCP Servers Report” on page 1147](#)
- ♦ [“Router Report” on page 1147](#)
- ♦ [“Segment Report” on page 1147](#)
- ♦ [“Segment Topology Report” on page 1147](#)

NOTE: The NCP Servers report is available only at the site level.

The Custom Atlas reports list information of nodes in the current container from which Reporting has been launched, and not of the nodes contained in the sub-level containers.

Prerequisites for Generating the Reports

Because Crystal Reports is invoked by DLLs on the system, you need to install the Sybase ODBC driver. To check if the driver is installed:

- 1** From the desktop *Start* menu, click *Settings > Control Panel > ODBC Data Source*.
 - 1a** In the *System Data Source Name (DSN)* pane, click *Add*.
 - 1b** Select the Adaptive Server Anywhere driver. You must install Adaptive Server Anywhere if you do not have it on your system. You can install it from the `sybase.zip` file at `companioncd\odbc\sybase*.*`
- 2** If you have an older version of Novell ZENworks Server Management, you will need to uninstall it and install the latest version of Novell ZENworks Server Management before you can run the reports.

To uninstall the previous version:

- 2a** From the desktop *Start* menu, click *Settings > Control Panel > Add/Remove Programs*.
- 2b** Select *Novell ConsoleOne* from the list and remove it.

If you have already installed the latest version, then delete the `zensnapins.jar` file from `novell consoleone\lib\zen`.
- 3** You will need at least MDAC 2.6 SP1 (Microsoft Data Access Component) for running Crystal Reports, particularly on a Windows machine. Check the version of MDAC on your box: select *Control panel > ODBC Data sources > the About* tab. The minimum version required is 3.520.7326.0. If the version you have does not match the minimum requirement, you need to upgrade the ODBC core components by downloading from [Microsoft site \(http://microsoft.com/data/download.htm\)](http://microsoft.com/data/download.htm).

Computer Systems by Segment Report

This report lists the number of computer systems on the selected segment. If the report is generated at the site level, the report lists the number of systems on each segment. For each segment, the report provides the following information about each connected computer system:

- ◆ Segment Name
- ◆ Segment Type
- ◆ Total nodes on a segment
- ◆ Node Name
- ◆ Node Address
- ◆ Services
- ◆ MIBs

NCP Servers Report

This report lists the following information for each server on the selected Novell ZENworks Server Management site:

- ◆ Server Name
- ◆ Total NCP servers on the site
- ◆ Server Label
- ◆ Server Address
- ◆ Labels (other names by which the server is known)
- ◆ MIBs

Router Report

This report provides the following information for each router on the selected Novell ZENworks Server Management segment or site:

- ◆ Total number of routers on the segment or site
- ◆ IPX Address
- ◆ Bound Segments
- ◆ Services
- ◆ MIBs
- ◆ IP Address
- ◆ MAC Address

Segment Report

This report lists the number of computer systems on the selected segment (segment level) or on all segments in the Novell ZENworks Server Management site (site level). For each segment, the report provides the following information about the systems connected to the segment:

- ◆ Segment Name
- ◆ Segment Type
- ◆ Total segments on the site
- ◆ IP configuration
- ◆ IPX configuration
- ◆ Total nodes on the segment

Segment Topology Report

This report provides information about the routers and bridges on a selected Novell ZENworks Server Management segment or site.

For each router, the report provides the following information:

- ◆ Router Name
- ◆ IP Address

- ◆ IPX Address
- ◆ MAC Address
- ◆ Bound Segment

For each bridge, the report provides the following information:

- ◆ Bridge Name
- ◆ Bridge Type
- ◆ Number of Ports
- ◆ Port: MAC Address and Bound Segment

33.1.2 About the Alarm Reports

The alarm reports provide information about the alarms received by the Novell ZENworks Server Management server. There are two types of alarm reports you can generate: Alarm details report and Alarm summary report.

This section provides information on the following topics:

- ◆ [“Prerequisites for Generating the Reports” on page 1148](#)
- ◆ [“Alarms Details Report” on page 1148](#)
- ◆ [“Alarms Summary Report” on page 1149](#)
- ◆ [“Available Trap Information Report on Site” on page 1149](#)

NOTE: The Custom Atlas reports the alarm information from the nodes in the current container only on which Reporting has been launched, and not of the alarms from the nodes contained in the sub-level containers.

Prerequisites for Generating the Reports

Because Crystal Reports is invoked by DLLs on the system, ensure that you have installed the Sybase ODBC driver. For more information, see [“Management and Monitoring Services Installation”](#) in the *Novell ZENworks 7 Server Management Installation Guide*.

Alarms Details Report

This report lists Information of the alarms on the site. The report is generated based on the customized settings. The report provides the following information about each connected computer system:

- ◆ Alarm Severity
- ◆ Affected object name
- ◆ Source address
- ◆ Alarm state
- ◆ Alarm category
- ◆ Alarm generator
- ◆ Alarm time

- ◆ Alarm owner
- ◆ Alarm type
- ◆ Alarm summary

Alarms Summary Report

This report generates a brief summary of the alarms on the site. It provides a graphical representation of the distribution of alarms, for the selected number of days. The report provides the following information about each connected computer system:

- ◆ Alarm Severity
- ◆ Alarm Category
- ◆ Alarm Owner
- ◆ Alarm state
- ◆ Top alarm types
- ◆ Top affected objects
- ◆ Top source address

Available Trap Information Report on Site

The Available Trap Information report lists the information of the SNMP traps currently available on the site server. The report is generated based on the MIBs compiled on the site server and provides the following information:

- ◆ Total traps
- ◆ Alarms category
- ◆ Alarm severity
- ◆ Alarm type
- ◆ Trap OID
- ◆ Trap description

33.1.3 About the Health Reports

The Health Reports provide information about the overall health of a specified Novell ZENworks Server Management site or network segment. Each health report is based on a predefined health profile. The health profiles define the trend parameters that are used to calculate the overall health of the segment or site. There are five predefined health profiles:

- ◆ [“Novell NetWare Server Profile” on page 1150](#)
- ◆ [“Microsoft Windows Profile” on page 1150](#)
- ◆ [“Ethernet Network Profile” on page 1150](#)
- ◆ [“Token Ring Network Profile” on page 1150](#)
- ◆ [“FDDI Network Profile” on page 1151](#)

In addition, you can modify any of the existing profiles or create your own health report profiles. See [“Customizing a Health Profile” on page 1152](#) or [“Adding a New Health Profile” on page 1152](#).

Novell NetWare Server Profile

Reports generated using this profile provide graphs of the following trend parameters and use these parameters to calculate the overall health of the Novell NetWare servers in the selected atlas, segment, or page:

- ◆ Cache Buffers
- ◆ Cache Hits
- ◆ CPU Utilization
- ◆ Volume Free Space

Microsoft Windows Profile

Reports generated using this profile use the following trend parameters to calculate health:

- ◆ Cache Hits
- ◆ CPU Utilization
- ◆ Disk Free Space
- ◆ Available Memory

In addition, reports generated using this profile contain trend graphs for the following parameter:

- ◆ Logged in Users

Ethernet Network Profile

Reports generated using this profile use the following trend parameters to calculate overall health:

- ◆ Total Errors
- ◆ Network Utilization

In addition, reports generated using this profile contain trend graphs for the following parameters:

- ◆ CRC error packets
- ◆ Undersized packets
- ◆ Oversized packets
- ◆ Fragmented packets
- ◆ Jabbers

Token Ring Network Profile

Reports generated using this profile use the following trend parameters to calculate overall health. In addition, reports generated using this profile also contain trend graphs for the following parameters:

- ◆ Network Utilization
- ◆ Total Errors

FDDI Network Profile

Reports generated using this profile use the following trend parameters to calculate overall health:

- ◆ Total Errors
- ◆ Network Utilization

In addition, reports generated using this profile contain trend graphs for the following parameters:

- ◆ CRC error packets
- ◆ Undersized packets
- ◆ Oversized packets
- ◆ Lost frame errors

33.2 Managing Reporting

The following sections provide procedures for customizing, generating, printing, and exporting the Novell ZENworks Server Management reports:

- ◆ [Section 33.2.1, “Managing the Topology Reports,” on page 1151](#)
- ◆ [Section 33.2.2, “Managing the Server Management Health Reports,” on page 1151](#)

33.2.1 Managing the Topology Reports

You can generate two types of topology reports: site-level reports and segment-level reports. The site-level reports provide details about the discovered devices on each segment in the Novell ZENworks Server Management site. The segment-level reports provide information about the managed devices on the selected network segment. Note that the NCP Servers report is available only at the site level.

The following section describes how to generate, print, and export a topology report.

Generating a Topology Report

To generate a topology report:

- 1 Select the Novell ZENworks Server Management site object, a network segment object, or a Custom Atlas container object.
- 2 Click *Tools > Reports*.
- 3 Select the report you want to generate, then click *Run Selected Report*.
- 4 To print the report, click .
- 5 To export the report, click .

33.2.2 Managing the Server Management Health Reports

The server management component provides five standard profiles that you can use to generate health reports. You can set up reports based on these standard profiles or you can customize these profiles or create your own profiles on which to base your reports. For information about the standard health profiles, see [“About the Health Reports” on page 1149](#).

This section contains the following tasks:

- ◆ [“Customizing a Health Profile” on page 1152](#)
- ◆ [“Adding a New Health Profile” on page 1152](#)
- ◆ [“Creating and Scheduling Health Reports” on page 1153](#)
- ◆ [“Editing, Scheduling, and Deleting Health Reports” on page 1153](#)
- ◆ [“Viewing and Printing a Health Report” on page 1154](#)
- ◆ [“Running a Health Report” on page 1155](#)
- ◆ [“Calculating the Overall Health” on page 1155](#)

Customizing a Health Profile

To customize a health profile:

- 1 Right-click the Novell ZENworks Server Management site object, then click *Properties*.
- 2 Select the *Health Profiles* tab.
- 3 Select the health profile you want to customize, then click *Edit*.
The Edit Profile dialog box is displayed. This dialog box contains a list of the parameters that can be used to calculate the overall health of the device or segment to which the profile is applied.
- 4 Specify the directory location to which reports generated using this profile should be published by entering a value in the Publish Directory field.
To browse for a directory, click the *Browse* button (...).
- 5 Modify the parameters that are used to calculate health by checking or unchecking the *In Health Calculation* check box next to each parameter. For more information on the parameters that are used in health calculation see, [“About the Health Reports” on page 1149](#).
- 6 Rank the importance of each trend parameter in calculating health by entering a number in the Weight field for each parameter you checked to include in the health calculation.
You can enter any whole number in the Weight field. The system will use the weights to determine how important the parameter is in calculating overall health. The larger the number, the more weight the parameter is given in calculating health.
- 7 Modify which parameters to render graphically in the health report by checking or unchecking the Show Trend on Report check box next to each parameter.
- 8 To save your changes, click *OK*.

Adding a New Health Profile

To add a new health profile:

- 1 Right-click the Novell ZENworks Server Management site object, then click *Properties*.
- 2 Select the *Health Profiles* tab.
- 3 Click *New*.
The New Profile dialog box is displayed.
- 4 Enter a name for the new profile in the *Name* field.

- 5 Select the type of device or segment to which the profile applies from the *Type* drop-down list then click *OK*.

The Edit Profile dialog box is displayed.

- 6 Specify the directory location to which reports generated using this profile should be published by entering a value in the *Publish Directory* field.

To browse for a directory, click the *Browse* button (...).

- 7 Select the parameters you want to use to calculate health for reports generated using this profile by clicking the *In Health Calculation* check box next to the appropriate parameters. For more information on the parameters that are used in health calculation see, [“About the Health Reports” on page 1149](#).

- 8 For each parameter you selected to include in the health calculation, indicate how important the parameter is in calculating overall health by entering a value in the *Weight* column.

You can enter any whole number in the *Weight* field. The system will use the weights to determine how important the parameter is in calculating overall health. The larger the number, the more weight the parameter is given in calculating health.

- 9 For each parameter that you want to be represented graphically in associated health reports, click the *Show Trend on Report* check box.

- 10 Click *OK*.

Creating and Scheduling Health Reports

To create and schedule a health report:

- 1 Right-click the container object, then click *Properties*.

- 2 Select the *Health Reports* tab.

- 3 Click *New*.

The Edit Report dialog box is displayed.

- 4 Enter a name for the report in the *Name* field.

- 5 Select the profile to use when generating the report by selecting a value from the *Profile* drop-down list.

- 6 Indicate how often you want to generate the reports by selecting a value from the *Period* drop-down list.

You can choose to generate reports daily, weekly, or monthly.

- 7 Set the time and date you want the reports generated by selecting or entering the appropriate values in the *Start Time*, *Day of the Week*, and/or *Day of the Month* fields.

The available fields will depend on the period you selected.

- 8 Click *OK*.

The report will be generated at the date and time you entered and stored in the directory specified in the associated report profile. For information on viewing the reports, see [“Viewing and Printing a Health Report” on page 1154](#).

Editing, Scheduling, and Deleting Health Reports

To edit and schedule a health report:

- 1 Right-click the atlas, page, or segment, then click *Properties*.

2 Select the *Health Reports* tab.

3 Click *Edit*.

The Edit Report dialog box is displayed. Edit the required information

4 Click *OK*.

IMPORTANT: If you want to edit the schedule time of the report, it is recommended that you create a new report with the changed schedule time or delete the report.

To delete a health report:

1 Right-click the atlas, segment, or page, then click *Properties*.

2 Select the *Health Reports* tab.

3 Click *Delete*.

4 Click *OK*.

Viewing and Printing a Health Report

After you create a health report, the report will be automatically generated on the day and time you specified. You can view the reports using a Web browser to open the `index.htm` file in the directory that is designated as the publish directory in the associated report profile.

IMPORTANT: Before you can view the health reports, you must install the Java plug-in 1.3.1_01 or higher version, except the 1.4.0 series. You can obtain this plug-in from Sun Microsystems, Inc.

To view a health report:

1 Browse to the directory where the health reports for the associated profile are stored.

2 Use your browser to open the `index.htm` file.

The `index.htm` file is a Java file containing all reports that are stored in the directory. The left column of the `index.htm` file lists report hierarchy.

3 Click the plus sign next to the profile that is associated with the reports you want to view.

The profile object expands to display a list of container objects.

4 Click the plus sign next to the container object associated with the reports you want to view.

The object expands to display a list of report names associated with the object.

5 Click the plus sign next to the report you want to view.

The object expands to display a list of individual report instances. For example, a report that is scheduled to run daily will have a report instance for each day. The reports are named by date and time. For example, `2000.09.09_11.15.10_PDT` is the name assigned to a report generated on September 9, 2000 at 11:15:10 Pacific daylight time.

6 Click the plus sign next to the report name to display a list of individual report pages.

The number of individual report pages depends on what report profile you selected and the object where you generated the report. For example, if you generated a report at the segment level using the Ethernet Network profile, there will only be one report page for the segment. If you generated a report at the site level using the Ethernet Network profile, there will be a report page for each Ethernet segment within the site. If you generated a report at the segment level

using the Novell NetWare Server profile, there will be a separate report page for each Novell NetWare server on the segment.

- 7 Click an individual report page to display the health report in the right frame.

The top of the report displays statistical information about the segment or server and provides a calculation of overall health. The parameters used to determine overall health are defined in the associated health report profile. The bottom of the profile displays trend graphs depicting the overall performance of the server or segment. See [“About the Health Reports” on page 1149](#) for a list of the parameters tracked and graphed in each of the standard profiles.

- 8 To print the report, click the *Print Report* button at the bottom of the left frame.

Running a Health Report

Although Health Reports are usually scheduled to run at a specified time of the day, week, or month, you may occasionally want to generate a Health Report on demand. To generate a Health Report on demand:

- 1 Right-click the atlas, segment, or page, then click *Properties*.
- 2 Select the *Health Reports* tab.
- 3 Select the report you want to generate, then click *Now*.

The report is saved to the directory specified in the report profile. See [“Viewing and Printing a Health Report” on page 1154](#).

Calculating the Overall Health

Overall health is calculated using the following parameters:

- ◆ Attributes selected for health calculation.
- ◆ Associated weights assigned to each attribute.
You can only associate weights, which are used for health calculations.
- ◆ Values for each attribute
Yellow threshold (YT), Red threshold (RT), and maximum value (maxValue).
- ◆ Global threshold values
Global Green threshold (GG) is 100, Global Yellow threshold (GY) is 66, and Global Red threshold (GR) is 33.

Health Calculation

For each of the attribute used in overall health calculation, sample values based on the schedule specified while generating the reports are collected. These sample values are normalized using global thresholds and attribute thresholds, where Global Green is 100, Global Yellow is 66, and Global Red is 33. The global Green range = global Green - global Yellow; the global Yellow range = global Yellow - global Red; and the global Red range = global Red.

Normalization Formula

Normalized Value = Global Threshold - ((value - attribute Threshold)/(attribute Threshold Range) * (Global Range)

if (value > attribute's RED threshold)

global Threshold = global Red
 attribute threshold = attribute Red threshold
 attribute threshold range = attribute max Value - attribute Red threshold
 global Range = global Red range
 if (value > attribute's Yellow threshold)
 global threshold = global Yellow
 attribute threshold = attribute Yellow threshold
 attribute threshold range = attribute Red - attribute Yellow
 global range = global Yellow range
 if (value > 0)
 global threshold - ((value)) / (attribute threshold range) * (global range)
 global threshold = global Green
 attribute threshold Range = attribute Yellow threshold
 global range = global Green range

Each of these may have an associated weight attached to it, which is configured in the respective profiles. Each of these attribute samples is then multiplied by the corresponding weights using the formula:

value = value * attributeWeight / TotalWeight;

where — value is the particular sample after normalization, attributeWeight is the weight associated with the attribute and the TotalWeight is the total weight of all the attributes used in health calculation.

The other values displayed in Health Reports are based on the following calculations:

- ◆ Minimum Value = minimum of all the values in a given sample
- ◆ Maximum Value = Maximum value of all the values in a given sample
- ◆ Average Value = Sum of all the Values / no of Samples
- ◆ Trend is calculated based on the Slope:

$$\text{Slope} = (n * \sum x * y - \sum x * \sum y) / (n * \sum x * x - \sum x * \sum x)$$

where:

n = number of samples

x = time at which these samples were captured

y = trend values

if Slope > 0, then the trend is increasing

if Slope < 0, then the trend is decreasing

if Slope = 0, then the trend is steady

- ◆ Intercept = ($\sum y$ - Slope * $\sum x$) / n

- ◆ Next Week Projection or Next Month Projection Value = Slope * time + Intercept
where time = Report Schedule Time (time when the report was scheduled) + 7 * 24 * 60 * 60 * 1000 for weekly Projection
- ◆ Report Schedule Time (time when the report was scheduled) + 30 * 24 * 60 * 60 * 1000 for Monthly Projection.

WARNING: Exporting data in CSV (Comma Separated Value), Character Separated Value, and Tab Separated Value (TSV), does not export the complete data. As a workaround you need to first export data in MS Excel format and then save it in the desired format.

If you export the generated reports in formats other than HTML or DHTML, the correct page numbers are not displayed. The page number is displayed incorrectly as Page -1 of 1, for all pages.

This chapter is referenced from the other sections. This section provides you information on SNMP, the SNMP community strings and how to configure SNMP community strings.

This section contains the following information:

- ♦ [Section 34.1, “About SNMP Community Strings,” on page 1159](#)
- ♦ [Section 34.2, “Setting the SNMP Community Strings,” on page 1160](#)

34.1 About SNMP Community Strings

SNMP is a protocol that offers network management services within the Internet suite of protocols.

SNMP uses a lightweight security mechanism whereby each protocol data unit (PDU) contains a community string. The SET community string is used in an SNMP Control operation and the GET community string is used in an SNMP Monitor operation.

SNMP community strings provide only a rudimentary form of security because they are transmitted in clear text in each SNMP request. Therefore, the community strings are exposed to any stations capable of monitoring an IP or Internetwork Packet Exchange™ (IPX™) network

Because Management Agent for Novell NetWare and Management Agent for Windows are based on SNMP, all actions that are directed from network Novell ConsoleOne to a server involve SNMP SET and GET requests from the manager to the agent. Novell ConsoleOne® requests data from a managed server by issuing an SNMP GET request. An SNMP SET command is required to set server alarm thresholds or configuration parameters. In most cases, you are unaware of the underlying SNMP commands required to carry out requests you make in Novell ConsoleOne, unless you are issuing requests on an SNMP-enabled device through the MIB Browser.

34.1.1 SNMP Security

Conducting management operations from Novell ConsoleOne raises the issue of ensuring security. In particular, if unauthorized users configuration parameters on a server, performance problems or even sabotage network operations are encountered.

For these reasons, you should establish a scheme for changing the default community string PUBLIC to a proprietary community string used for communication between the management system and your SNMP agents.

Use the community keyword to define the community string to be used in the generated traps. The length of the community string is restricted to 32 bytes and cannot contain a space (except between quotes), tab, square bracket, equals sign, colon, semicolon, or number sign (#) characters. You can use Unicode* or International characters for the community string.

The default community string for Monitor operations is PUBLIC and for Control operations is null.

34.2 Setting the SNMP Community Strings

This section provides the following information:

- ◆ [Section 34.2.1, “Setting the SNMP Community String: Novell NetWare Server,”](#) on page 1160
- ◆ [Section 34.2.2, “Setting the SNMP Community String: Novell ConsoleOne,”](#) on page 1162
- ◆ [Section 34.2.3, “Setting Community Strings for an Individual Node,”](#) on page 1162
- ◆ [Section 34.2.4, “Setting the SNMP Community String: Windows,”](#) on page 1163

34.2.1 Setting the SNMP Community String: Novell NetWare Server

You configure security access for SNMP communications using either SNMP LOAD command line parameters (Novell NetWare 3.x/4.x/5.x/6 servers) or through INETCFG (Novell NetWare 4.x/5.x/6 servers, or servers with Novell NetWare MultiProtocol Router™ software installed).

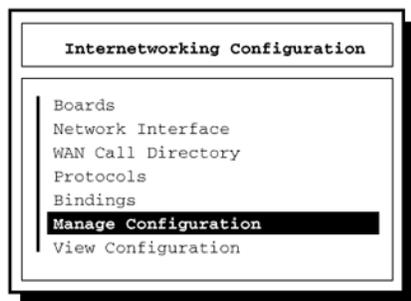
The following sections contain additional information to help you configure your Novell NetWare servers:

- ◆ [“Configuring Community String Options Using INETCFG”](#) on page 1160
- ◆ [“Configuring Community String Options Using SNMP LOAD Commands”](#) on page 1161

Configuring Community String Options Using INETCFG

To configure the community string options using INETCFG:

- 1 At the server prompt, enter `LOAD INETCFG`.



- 2 From the Internetworking Configuration menu, click *Manage Configuration > Configure SNMP Parameters > Monitor State*.
- 3 Select one of the following options:

These options let you indicate how SNMP handles SNMP read operations coming from outside this server.

| Option | Description |
|-------------------------------|--|
| <i>Any Community May Read</i> | Allows all GET (read) commands no matter what community string is provided in the incoming read request. |

| Option | Description |
|-------------------------------------|---|
| <i>Leave as Default Setting</i> | Avoids changing the Monitor community string from its default (which is usually PUBLIC). The default Monitor Community can still be changed manually through SNMP command line options, as described in “Configuring Community String Options Using SNMP LOAD Commands” on page 1161 . |
| <i>No Community May Read</i> | Allows GET (read) commands only for requests that are made by Novell ConsoleOne that have logged in to the server with SUPERVISOR or OPERATOR privileges. Any community string provided in an incoming read request is ignored. |
| <i>Specified Community May Read</i> | Allows only GET (read) commands for requests that contain the name specified in the Monitor Community field. If you selected this option, enter a name in the Monitor Community field, then press Enter. Enter the name of the community that is allowed to read management information. SNMP management stations that belong to this community can read the network management database. |

4 Press Enter.

To change the Control community options, repeat Step 1 to Step 4 and choose the appropriate options for the community strings.

5 When you are finished, press Esc. If prompted, click *Yes* to save changes to the SNMP parameters, then press Enter.

6 To return to the Internetworking Configuration menu, press Esc.

7 To exit INETCFG, press Esc.

8 Re-initialize the system.

To re-initialize, at the server prompt, enter `reinitialize system`.

Configuring Community String Options Using SNMP LOAD Commands

The LOAD command accepts the following SNMP option parameters:

- ♦ **MonitorCommunity:** Sets the community string for read-only (GET) access. The default value is PUBLIC. The syntax is as follows:
`LOAD SNMP MonitorCommunity=community_name`
- ♦ **ControlCommunity:** Sets the community string for read and write (GET and SET) access. By default, this community string is disabled.

The syntax is as follows:

`LOAD SNMP ControlCommunity=community_name`

These options set the community string for the indicated community.

The following table shows examples of available settings:

IMPORTANT: Community strings are case sensitive.

| Access Available to Requester | Read Only | Read/Write |
|-------------------------------------|--|--|
| Community name: "secret" | Load SNMP MonitorCommunity= <i>secret</i> | LOAD SNMP ControlCommunity= <i>secret</i> |
| | or | |
| | LOAD SNMP ControlCommunity= <i>secret</i> | |
| Community name: "str1" or "str2" | Load SNMP MonitorCommunity= <i>str1</i> | |
| | and | |
| | LOAD SNMP ControlCommunity= <i>str2</i> | |
| Any community name | Load SNMP MonitorCommunity="" | LOAD SNMP ControlCommunity="" |
| | or | |
| | LOAD SNMP ControlCommunity="" | |

34.2.2 Setting the SNMP Community String: Novell ConsoleOne

You set global community and trap target information using the SNMP property page associated with the site-level object. You can also customize the setting for a specific device using the SNMP property page of the device itself.

34.2.3 Setting Community Strings for an Individual Node

This section describes the procedure to set up the community strings for SNMP SET and GET operations on an individual node.

Typically, community strings are configured to be identical over all nodes in a network, or at least over a portion of the network. The default value for both SET and GET is public. The community strings are case sensitive.

By default, Novell ZENworks Server Management uses the public community string for SNMP GET and SET operations. You can configure a community string other than public on a node-by-node basis, or you can configure a community string globally on all SNMP-managed nodes. The community string that Novell ZENworks Server Management uses must match the string expected by the SNMP agent in the managed node; otherwise, the operation will fail.

To set up the community strings for SET and GET operations for an individual node:

- 1 In Novell ConsoleOne, click the target SNMP-manageable node.
- 2 Right click the node, then click *SNMP Settings*.
- 3 Enter the community string.

Novell ZENworks Server Management uses this community string for SET and GET operations when communicating with the device.

4 Click *OK*.

34.2.4 Setting the SNMP Community String: Windows

You configure security access for SNMP communications on Windows servers using the Network applet in the Windows Control Panel. For detailed information, refer to your Windows documentation or online help.

You must load the Microsoft* SNMP Service on your Windows servers. The SNMP community string setting must be the same as the SNMP community string setting on your Novell ConsoleOne.

Understanding the View Builder

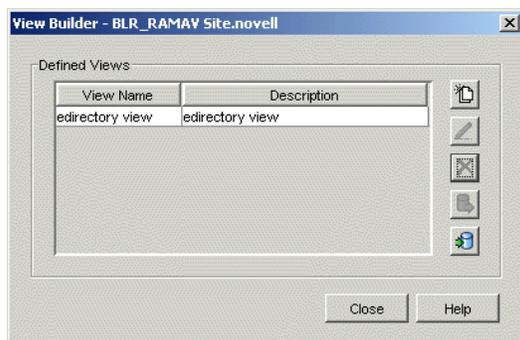
35

The View Builder provides a mechanism through which you can create a view to display information from the agents that have instrumented SNMP MIBs and traps sent by the agent to the Management Site Server. You can use the View Builder to create views in addition to those available in Novell ConsoleOne. These views are displayed as text, tables and graphs.

You can associate the views to specific nodes and manage them.

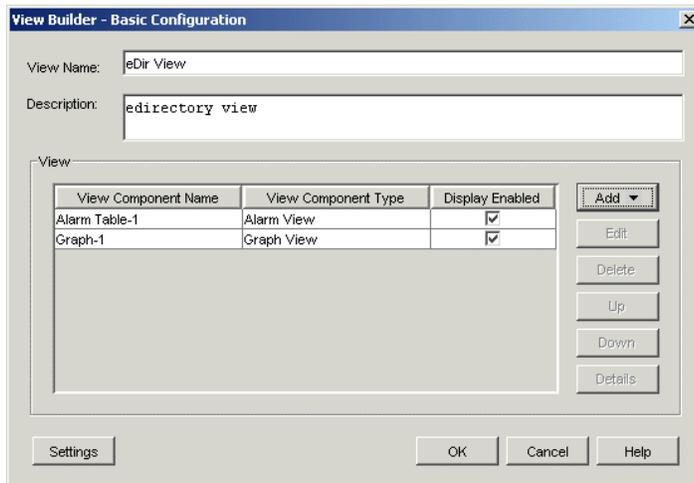
To launch the View Builder dialog box:

- 1 In Novell ConsoleOne, select *Tools > View Builder*.



- ♦ **Defined Views:** Lists the views created.
 - ♦ **View Name:** The name of the view.
 - ♦ **Description:** The description of the view.
 - ♦ **Add:** Click . Launches the Basic Configuration dialog box. You can create the view components in this dialog box.
 - ♦ **Edit:** Click  to edit the view details.
 - ♦ **Delete:** Click  to delete the view.
 - ♦ **Export:** Click  to export a view to a file.
 - ♦ **Import:** Click  to import a view from a file.
- 2 In the View Builder dialog box, click .

The View Builder - Basic Configuration dialog box is displayed.



A view can contain view components or a combination of view components. The view components you can create are: Name-Value Pairs view, Alarm View, Table View, and Graph View.

3 Perform any of the following operations:

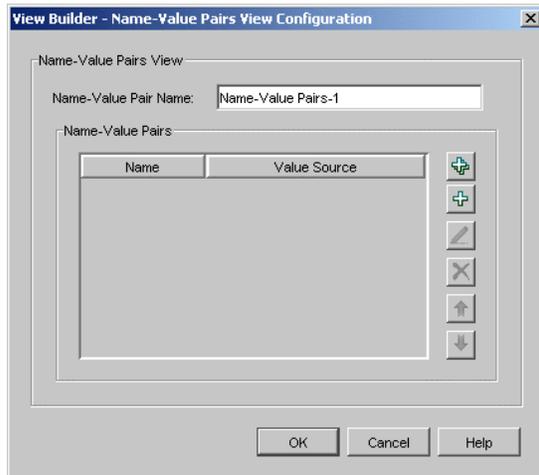
- ◆ [Section 35.1, “Creating a Name-Value Pairs View Component,” on page 1167](#)
- ◆ [Section 35.2, “Creating an Alarm View Component,” on page 1169](#)
- ◆ [Section 35.3, “Creating a Table View Component,” on page 1170](#)
- ◆ [Section 35.4, “Creating a Graph View Component,” on page 1171](#)
- ◆ [Section 35.5, “Setting the Criteria for the View to Appear,” on page 1172](#)

4 Click *OK*.

35.1 Creating a Name-Value Pairs View Component

The Name-Value Pairs view component can consist of multiple name-value pairs.

- 1 In the View Builder - Basic Configuration dialog box, click *Add > Name-Value Pairs View*.



- 2 Specify the *Name-Value Pairs View* name.

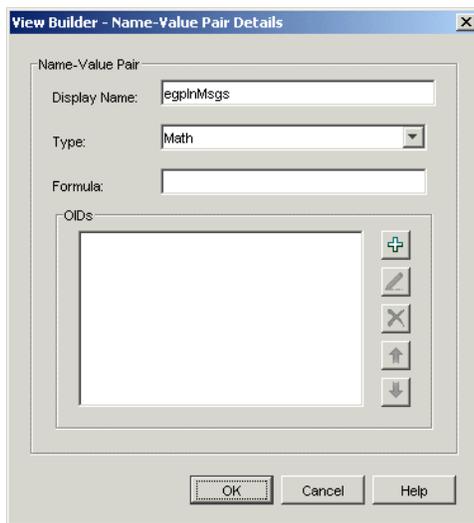
The *Name-Value Pairs* box contains a tabular list of all the view components you have created. The Name and the Value Source of the view component you created are displayed.

- 3 Perform any of the following operations:

- ♦ **Add Multiple:** To add multiple name-value pairs, click .
- ♦ **Add:** To add a name-value pair, click .
- ♦ **Edit:** To modify the information, select the name-value pair you want to edit from the Name-Value Pairs list, then click .
- ♦ **Delete:** To delete a name-value pair, select it from the Name-Value Pairs list, then click .
- ♦ **Arrange Order:** Use  or  to change the order of display.

35.1.1 Adding a Name-Value Pair

- 1 In the Name-Value Pairs View Configuration dialog box, click .



- 2 Specify the display name of the name-value pairs component.
- 3 From the drop-down list, select the type of the name-value pair.

You can select *Name-Value Pair* type as *Simple*, *Rate*, or *Math*. If you select the type as *Rate*, you must specify OIDs of types GAUGE or COUNTER. If you select the type as *Math*, you must specify a formula.

Specifying a Formula: The OIDs listed in the table are numbered from 0,1,2 and so on. They are preceded by the data identifiers v, the current value of the OID; and h, indicating the previous value of the OID. You can use +, -, /, *, ^, () operands when you specify your formula. You can also use constants in your formula.

For example: Select the OIDs snmpInPkts (1.3.6.1.2.1.11.1) and snmpOutPkts (1.3.6.1.2.1.11.2). The formula can be $v0+v1$. The result will be the total number of SNMP packets handled by the transport service.

- 4 Click *OK*.

The name-value pair you created is displayed in Name-Value Pairs Configuration dialog box. Use the same procedure to add multiple name-value pairs, then Click *OK*. The Name-Value Pair view component is displayed in the Basic Configuration dialog box.

35.1.2 Editing the Name-Value Pairs View Component

- 1 In the Basic Configuration dialog box, select the name-value pairs view component you want to modify, then click *Edit*.
- 2 Modify the name-value pairs component details, then click *OK*.

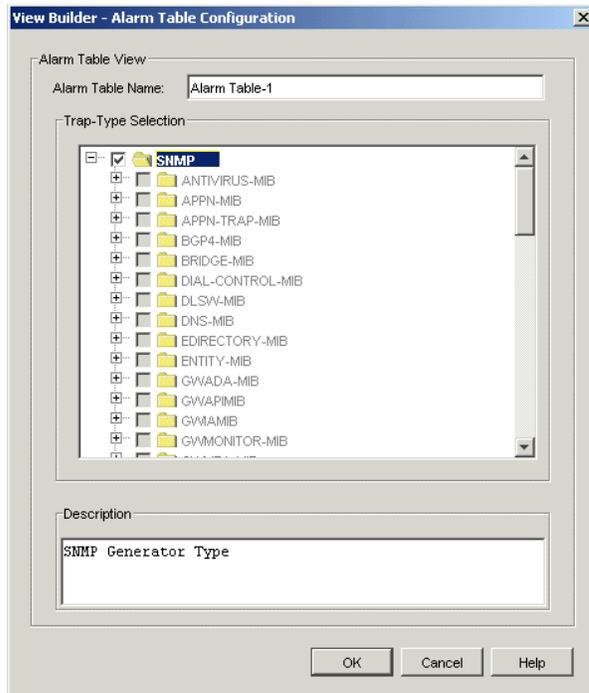
35.1.3 Deleting the Name-Value Pairs View

- 1 In the Basic Configuration dialog box, select the name-value pairs view component you want to delete.

- 2 Click *Delete*.

35.2 Creating an Alarm View Component

- 1 In the View Builder - Basic Configuration dialog box, click *Add > Alarms View*.



- 2 Specify the Alarm Table name.
- 3 Select a trap type.

You can select the trap type based on three hierarchies. You can select all the traps from the node, a particular category, or a specific trap type within a category.

- 4 Click *OK*.

The alarm view component you created is displayed in the Basic Configuration dialog box.

35.2.1 Editing the Alarms View Component

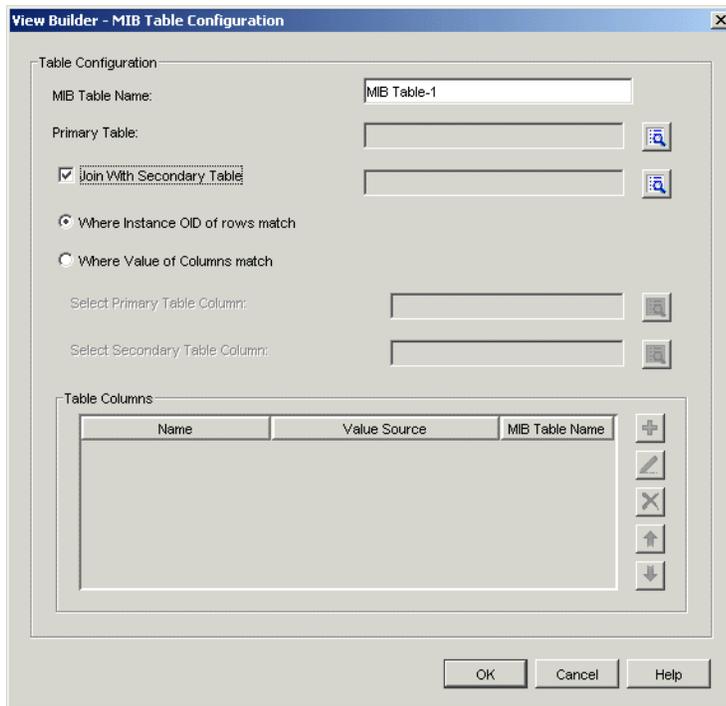
- 1 In the Basic Configuration dialog box, select the alarm view component you want to modify, then click *Edit*.
- 2 Modify the alarm view component details, then click *OK*.

35.2.2 Deleting the Alarm View Component

- 1 In the Basic Configuration dialog box, select the alarm view component you want to delete.
- 2 Click *Delete*.

35.3 Creating a Table View Component

- 1 In the View Builder - Basic Configuration dialog box, click *Add > Table View*.



- 2 Specify the MIB Table name.
- 3 Click  to select the MIB table that will be your primary table. Go to [Step 7](#).
- 4 Join the primary table with the secondary table.
- 5 Check the *Join With Secondary table* option if you want to join the primary and secondary table. Click  to select the secondary table.
- 6 Set the conditions to join the primary and secondary table

Where Instance OID of Rows Match: Joins the primary and the secondary table based on matching instance OIDs.

Often extension tables are created for standard tables. The extended table does not have an explicit index and uses the index of the standard table. For example, you can join the tables `hrSWRunTable(1.3.6.1.2.1.25.4.2)` and `hrSWRunPerfTable (1.3.6.1.2.1.25.5.1)`.

Where Values of Columns Match: Joins the primary and the secondary table based on matching columns OIDs. You must select the primary table column and the secondary table column to meet this condition. Click  to select the Primary table column and the second table column.

For example, you can join the `ifTable(1.3.6.1.2.1.2.2)` and `ipAddrTable(1.3.6.1.2.1.4.20)` with the matching columns being `ifIndex:1.3.6.1.2.1.2.2.1.1` from primary table and `ipAdEntIfIndex 1.3.6.1.2.1.4.20.1.2` from the secondary table.

- 7 Under Table Columns perform any of the following operations:
 - ♦ **Add:** To add a table column, click .
 - ♦ **Edit:** To change the table column, click .

- ♦ **Delete:** To delete the table column, click .
- ♦ **Arrange Order:** Use  or  to change the order of display.

8 Click *OK*.

The table view component you created is displayed in the Basic Configuration View dialog box.

35.3.1 Editing an Table View

- 1 In the Basic Configuration dialog box, select the table view component you want to modify, then click *Edit*.
- 2 Modify the table view details, then click *OK*.

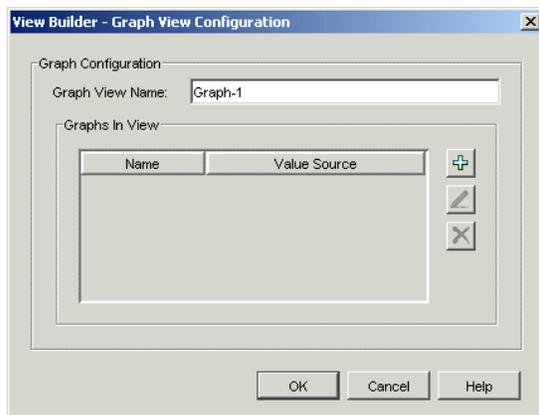
35.3.2 Deleting a Table View

- 1 In the Basic Configuration dialog box, select the table view component you want to delete, then click *Delete*.

35.4 Creating a Graph View Component

The graphs you want to add to the graph view are displayed based on the trend data collected by the Advanced Trending Agent running on a particular node. Use the following procedure to create graphs:

- 1 In the View Builder - Basic Configuration dialog box, click *Add*, then click *Graph View*.



- 2 Specify the graph view name.

35.4.1 Adding Graph Details

- 1 In the Graph View Configuration dialog box, click .



- 2 Specify the display name of the graph.
- 3 The Advanced Trending Agent is selected because this is the trending agent the graph display is based on.
- 4 Click  to select the trend OID.

IMPORTANT: The OID you select must be trended by the Advanced Trending Agent.

- 5 Click *OK*

The graph you created will be displayed in the Graph View Configuration dialog box. Use the above procedure to add multiple graphs, then click *OK*. The Graph view component is displayed in the Basic Configuration dialog box.

35.4.2 Editing the Graph View Component

- 1 In the Basic Configuration dialog box, select the graph view component you want to modify, then click *Edit*.
- 2 Modify the graph details, then click *OK*.

35.4.3 Deleting the Graph View Component.

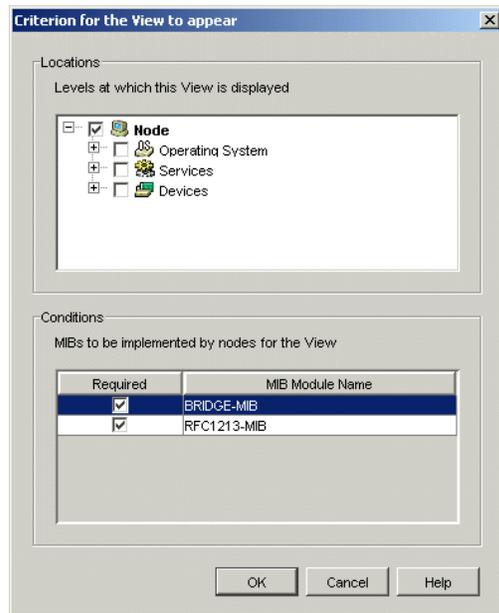
- 1 In the Basic Configuration dialog box, select the graph view component you want to delete.
- 2 Click *Delete*.

35.5 Setting the Criteria for the View to Appear

You can set a criterion for the view to appear on selected levels. The setting includes the location of the view and the condition for the view to appear.

To set the criterion for the view to appear:

- 1 In the Basic Configuration dialog box, click *Settings*.



- 2 Specify the location where you want the view to appear.
By default, the view you created appears at the node level.
- 3 Specify the condition for the view to appear.

The MIBs implemented by the node and the MIB Module name of the node where you want the view to be displayed are displayed in this list. By default, all the MIB modules are deselected for implementing on the node. If it is necessary for the MIB module to be implemented on a node, select the MIB module. The MIBs listed here will be automatically discovered by Discovery.

- 4 Click *OK*.

Novell® ZENworks® 7 Server Management provides the Trap Configuration tool to manage NetWare® Server Traps and Novell Directory Services® Traps in your network. Trap Configuration allows you to configure the Traps on the NetWare servers from a centralized management console.

The following sections provide you with the concepts and instructions to help you configure the Traps:

- ♦ [Section 36.1, “Understanding Trap Configuration,” on page 1175](#)
- ♦ [Section 36.2, “Configuring Traps Using Trap Configuration Page,” on page 1176](#)
- ♦ [Section 36.3, “Additional Trap Configuration Features,” on page 1179](#)

36.1 Understanding Trap Configuration

This section contains basic information to help you understand the ZENworks 7 Server Management Trap Configuration feature.

- ♦ [Section 36.1.1, “The Configuration Agents,” on page 1175](#)
- ♦ [Section 36.1.2, “Trap Configuration Management Console,” on page 1175](#)

36.1.1 The Configuration Agents

Trap Configuration allows you to configure Traps of NetWare Server Alarm MIB and Novell Directory Services MIB.

NetWare Server Alarm MIB Trap Configuration is managed by the `nwtrpagt.nlm` agent, which implements the NWTRAPCONFIGURATION MIB module.

The Novell Directory Services MIB Trap Configuration is managed by the `dstrpagt.nlm` agent, which implements the NDSTRAPCONFIGURATION MIB module.

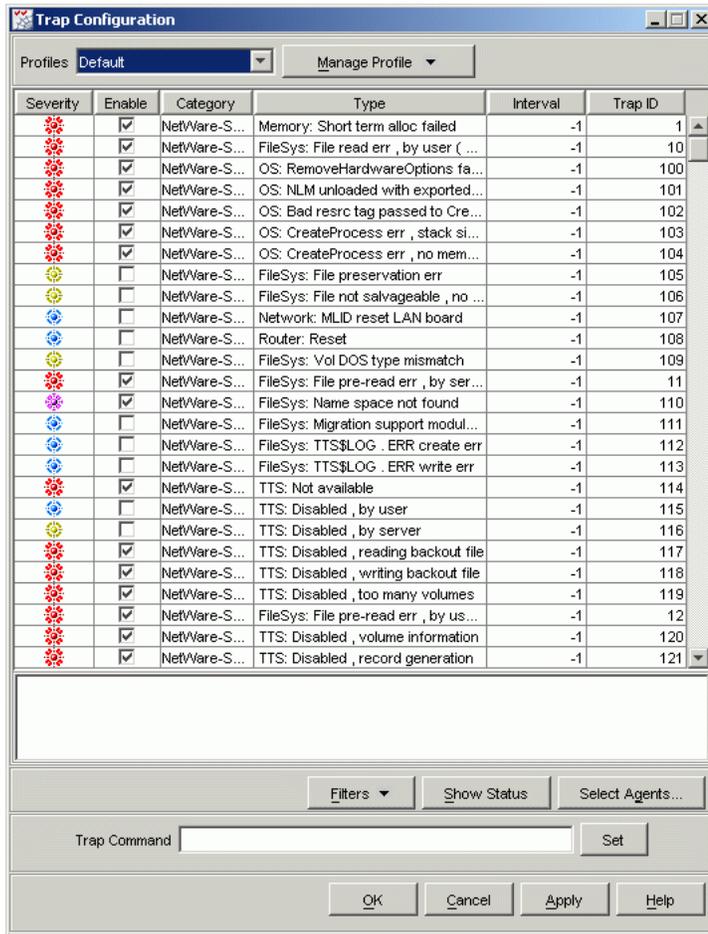
The `nwtrpagt.nlm` and `dstrpagt.nlm` agents implement a table for all the NetWare Server Alarm MIB Traps and Novell Directory Services MIB Traps. These tables are defined by NWTRAPCONFIGURATION MIB and NDSTRAPCONFIGURATION MIB. Each row in this table has an editable fields for Trap Status and Intervals. The number of rows are same as the number of Traps defined in NetWare Server Alarm MIB and Novell Directory Services MIB.

NOTE: You cannot create a row in this table.

36.1.2 Trap Configuration Management Console

You can use the Trap configuration management console to enable/disable and configure the Interval field for NetWare Server Traps and Novell Directory Service Traps, as illustrated in [Figure 36-1](#):

Figure 36-1 Trap Configuration dialog box



36.2 Configuring Traps Using Trap Configuration Page

Following are the tasks you can perform on the Trap Configuration Management Console using Trap Configuration page:

- ◆ [Section 36.2.1, “Enabling and Disabling the Traps,” on page 1176](#)
- ◆ [Section 36.2.2, “Changing the Interval of a Trap,” on page 1177](#)
- ◆ [Section 36.2.3, “Selecting NetWare Servers to Apply a Trap Configuration,” on page 1177](#)
- ◆ [Section 36.2.4, “Viewing the Trap Configuration Status,” on page 1178](#)
- ◆ [Section 36.2.5, “Using the Command Line Option,” on page 1178](#)

36.2.1 Enabling and Disabling the Traps

The Trap Configuration feature allows you to enable/disable Trap or a set of Traps. If you are accessing the Trap configuration of a NetWare server, the current configuration of a NetWare server is displayed; however if you are accessing the segment or atlas level, the default profile is displayed where Critical and Major Traps are enabled and Minor and Informational Traps are disabled.

To enable/disable the Traps:

- 1 In ConsoleOne, right-click the object under a Site Server object, then click *Actions*, then click *Trap Configuration*.
- 2 Select an option next to a Trap which you want to enable. When you select a row, the corresponding description is displayed in the bottom of the Trap Configuration dialog box.
You disable a Trap by deselecting the option.

NOTE: You must select the NetWare server where you want to enable or disable a Trap. For more details on selecting a NetWare server, refer to [Section 36.2.3, “Selecting NetWare Servers to Apply a Trap Configuration,” on page 1177](#)

36.2.2 Changing the Interval of a Trap

Trap Configuration allows you to change the Trap interval. An interval indicates the duration before the same Trap is generated again. Controlling the interval allows you to eliminate excess server traffic.

- 1 In ConsoleOne, right-click the object under a Site Server object, then click *Actions*, then click *Trap Configuration*.
- 2 Select a Trap for which you want to change the interval.
- 3 Double-click *Interval*.
- 4 Change the interval value according to your requirements.

NOTE: You can also change the Trap interval by using command line options. For more details on using command line options, refer to [Section 36.2.5, “Using the Command Line Option,” on page 1178](#)

36.2.3 Selecting NetWare Servers to Apply a Trap Configuration

You need to select a NetWare server or a set of NetWare servers to apply a configuration on it.

- 1 In ConsoleOne, right-click the object under a Site Server object, then click *Actions*, then click *Trap Configuration*.
- 2 Click *Select Agents*.
- 3 In Select Node View dialog box, click *Add*.
- 4 In Select Object dialog box, select the NetWare servers on which you want to apply the configuration.

NOTE: You can select more than one NetWare server by pressing Ctrl and selecting the servers you want to include.

To remove NetWare server, select the agent from the Select Node View dialog box, then click *Remove*.

- 5 Click *OK*. The IP address of the selected NetWare server is displayed in the Select Node View dialog box.
- 6 Click *OK*, then click *Apply*.

The selected Trap configuration is applied on the listed NetWare servers.

36.2.4 Viewing the Trap Configuration Status

- 1 In ConsoleOne, right-click the object under a Site Server object, then click *Actions*, then click *Trap Configuration*.
- 2 Click *Show Status*.

The status of previous configuration operation is displayed in the Configuration Status dialog box.

NOTE: To refresh the status, click *Refresh* or Press F5. If you have selected multiple servers, Trap Configuration might take some time to refresh the status.

36.2.5 Using the Command Line Option

Trap Configuration allows you to specify a command to:

- ◆ Enable/disable the Traps
- ◆ Set up an interval for the Traps

You can use various operators like =, >, <, >=, <= to work with Trap commands. For example, if you want to enable all NetWare server Traps with an ID of more than 20, type following in the Trap Command field:

```
nwTrap enable ID>20
```

The following table lists and explains the text or string you can specify in the *Trap Command* field:

| Text or String | Result |
|-----------------------------------|---|
| nwtrap enable all | Enables all NetWare traps. |
| nwtrap enable ID > 20 | Enables all NetWare traps with an ID greater than 20. |
| nwtrap enable severity < critical | Enables those NetWare traps that have a severity that is less than Critical. You can also use numbers to represent a severity. For example, you can use 4 for Critical and 3 for Major. |
| nwtrap all interval 20 | Sets the interval as 20 for all NetWare traps. |
| nwtrap 1-20 interval 10 | Sets the interval as 10 for the first 20 traps. |
| nwtrap 1 2 3 interval 15 | Sets the interval as 15 for the first three NetWare traps. |

NOTE: In the above table, nwtrap stands for NetWare server Traps. You can use similar commands for Novell Directory Services Traps by using ndstrap instead of nwtrap. For example, ndstrap enable all enables all Novell Directory Services Traps.

36.3 Additional Trap Configuration Features

ZENworks Server Management Trap Configuration also includes the following features:

- ♦ [Section 36.3.1, “Filtering the Traps,” on page 1179](#)
- ♦ [Section 36.3.2, “Sorting the Traps,” on page 1180](#)
- ♦ [Section 36.3.3, “Managing Profiles,” on page 1180](#)
- ♦ [Section 36.3.4, “Viewing Current Configuration of a Server,” on page 1181](#)
- ♦ [Section 36.3.5, “A Use Case: Configuring NetWare Servers,” on page 1181](#)

36.3.1 Filtering the Traps

You can display the Traps in the table based on filter conditions.

You set up a filter by selecting criteria from four drop-down lists. You can set up simple filters that contains only one line, or complex filters consisting of multiple lines or groups. If you set up a filter using more than one line, you must also specify the logical relationship between the lines.

To set up a filter:

- 1 In ConsoleOne, right-click the object under a Site Server object, then click *Actions*, then click *Trap Configuration*.
- 2 In Trap Configuration page, click *Filters*, then click the *Filter* menu.
- 3 In the first drop-down list, select the column by which you want to filter the Traps.
- 4 Select an operator from the second drop-down list.

The operator defines the constraint value set for the column. You can specify equal to, not equal to, less than, greater than, greater than or equal to, less than or equal to, contains, or start with.
- 5 Select a value from the third drop-down list.

The list of available values depends on the value you have selected in first drop-down list. For example, if you have selected Severity in the first drop-down list, the values in the third drop-down list are Critical, Major, Minor, Informational, and Unknown. If you have selected Enable in the first drop-down list, the available values are Enabled and Disabled.
- 6 Select a value from the fourth drop-down list to specify how this filter statement relates to other statements you plan to define.
 - ♦ If this is the only filter statement or if it is the last statement in a group, select End.
 - ♦ If you want to add a line below the current filter statement, select New Row. You must define the logical relationship between the previous line and the new line. The Traps are displayed based on the logical condition you have specified. Select And to satisfy both the filter conditions. Select Or to satisfy any one of the filter conditions.
 - ♦ If you want to add one or more lines that are unrelated to the preceding lines, select New Group. An additional drop-down list separates the new line from the preceding lines. Select a value from this drop-down list to indicate the relationship between the filter statements. Select And if you want both the filter statements to be satisfied. Select Or if you want only one of the filter statements in one of the groups to be satisfied. Select End from the fourth drop-down list when you add a new group.
- 7 Click OK.

The Trap configuration list displays only those entries that meet the filter criteria.

To clear the list and go back to the default view, click *Clear Filter* under *Filters*.

36.3.2 Sorting the Traps

You can modify the order in which the Traps are displayed in the Trap Configuration by sorting the Traps.

Double-click the column header for the column by which you want to sort the entries to sort the entries in the view in descending order. To sort the entries by ascending order, double-click the column header again.

36.3.3 Managing Profiles

Profile is a set of predefined configuration of the Traps. You can have a profile consisting of specific set of Traps. For example, you can create a profile only with File Sys Traps.

The Trap Configuration feature allows you to configure the different profiles and provides a central location for creating, managing, and controlling profiles. You can add a Trap or set of Traps to a profile.

The different tasks you can perform with Trap Configuration profile include:

- ◆ “Creating a New Profile” on page 1180
- ◆ “Editing a Profile” on page 1181
- ◆ “Deleting a Profile” on page 1181

Creating a New Profile

- 1 In ConsoleOne, right-click any object under a Site Server object, click *Actions*, then *Trap Configuration*.
- 2 In *Manage Profile* field, click the down-arrow, then click *New*.
- 3 On the New Profile page, click *Add* to open the Default Configuration View.

The Default Configuration View, where all Major and Critical alarms are enabled and interval is set to -1 for all the Traps, lists all NetWare server alarms and Novell Directory Services Traps.

- 4 Select the Traps to be included in the profile, then click *Add*.

NOTE: You can select more than one Trap by pressing Ctrl and selecting the Traps you want to include in the profile.

You can also use filters to display a specific set of Traps. For more information on filters, refer to [Section 36.3.1, “Filtering the Traps,” on page 1179](#).

- 5 In Manage Profile, click *Save* or *Save As*.
- 6 In the Profile List View, specify name of the profile in *Profile Name* text box, then click *Save*.

NOTE: If you use an existing profile name, the system asks you if you want to overwrite the existing profiles. If you click *Yes*, the existing profile is over-written by the profile you are creating.

The newly created profile is added to the *Profiles* drop-down list.

Editing a Profile

- 1 In ConsoleOne, right-click the node under a Site Server object, click *Actions*, then click *Trap Configuration*.
- 2 Select the profile you want to edit. The corresponding Traps are displayed in the Trap Configuration page.
- 3 To add more Traps to the profile list, click *Add* and add the desired Traps from Default Configuration View.

NOTE: You can also remove the Traps from the list by clicking on *Remove*. You can remove more than one Trap by pressing Ctrl and selecting the Traps you want to remove.

- 4 In *Manage Profile* field, click *Save*.
- 5 (Optional) To save the updated profile with different name, click *Save As*, specify the name in the *Profile Name* field, then click *Save*.

Deleting a Profile

- 1 In ConsoleOne, right-click Atlas, then click *Actions*, then click *Trap Configuration*.
- 2 In the *Profiles* drop-down list, select the profile you want to delete.
- 3 In *Manage Profile* field, click the down-arrow, then click *Delete*.
- 4 Click *Yes* in the Confirm Delete dialog box. The selected Profile is deleted from the drop-down list.

NOTE: You cannot delete the default profile.

36.3.4 Viewing Current Configuration of a Server

- 1 In ConsoleOne, browse through the tree to find the node for which you want to view the configuration.

You can also find the node by clicking *Find* menu in ConsoleOne.

- 2 Right click the node and select *Action*.
- 3 Select *Trap Configuration*.

The current configuration that is applied on the selected server is displayed.

36.3.5 A Use Case: Configuring NetWare Servers

Trap Configuration also allows you to configure specific set of Traps.

For example, you can use the following procedure to configure File Sys Traps:

- 1 In ConsoleOne, right-click the node under a site-server object, click *Actions*, then *Trap Configuration*.
- 2 Click *Filters*, and then click *Filter*.
- 3 Select *Type* in first drop-down list.
- 4 Select *contains* in the second drop-down list.
- 5 Specify *File Sys* in third drop-down list.

- 6 Select *End* in fourth drop-down list, then click *OK*.
A list of all File Sys type Traps is displayed.

Setting up Security for Management and Monitoring Services



Management and Monitoring Services should be secured to ensure protection of all components and the database.

- ♦ The software is designed to work in a secured network, behind a firewall. Make sure all the components are within a secured network or firewall.
- ♦ The database contains network topology information, that is vulnerable to hacking. Make sure the database is protected.
- ♦ MMS uses SNMP v1, which does not support any protection, because the community name is transferred in clear-text format. For SNMP communication to NetWare servers, there is an option to use SNMP over NCP, which ensures that authentication does not use any community names, but still has valid NCP connections. SNMP over NCP is not supported for Windows.
- ♦ Log files might contain information about the database user name, passwords, and community string if the debug option is set to a high level. Use high debug levels only when necessary to assist in debugging, and only for a limited time.

Documentation Updates

R

This section contains information on documentation content changes that have been made in the *Administration* guide for Management and Monitoring Services since the initial release of Novell® ZENworks® 7 Server Management. The information will help you to keep current on updates to the documentation.

All changes that are noted in this section were also made in the documentation. The documentation is provided on the Web in two formats: HTML and PDF. The HTML and PDF documentation are both kept up-to-date with the documentation changes listed in this section.

The documentation update information is grouped according to the date the changes were published. Within a dated section, the changes are alphabetically listed by the names of the main table of contents sections for Management and Monitoring Services.

If you need to know whether a copy of the PDF documentation you are using is the most recent, the PDF document contains the date it was published on the front title page or in the Legal Notices section immediately following the title page.

The documentation was updated on the following date:

- ♦ Section R.1, “August 16, 2006,” on page 1185
- ♦ Section R.2, “July 14, 2006 (Support Pack 1),” on page 1185
- ♦ Section R.3, “December 9, 2005,” on page 1186

R.1 August 16, 2006

Updates were made to the following sections:

| Location | Change |
|--|--|
| Section 30.3.3, “Specifying Negative Filter Conditions in the Nttrap.ini File,” on page 1122 | Added the following paragraph after the steps: For more information about filter conditions and examples in the <code>nttrap.ini</code> file, see TID 10098619 (http://support.novell.com/cgi-bin/search/searchtid.cgi?10098619.htm) in the Novell Support Knowledgebase. |

R.2 July 14, 2006 (Support Pack 1)

Updates were made to the following sections:

| Location | Change |
|--|------------------------|
| Section 23.3.2, “VLAN Atlas,” on page 915 | This section is added. |
| Chapter 27, “Using the Probe Manageability Tool,” on page 1019 | This section is added. |

| Location | Change |
|--|---|
| Appendix Q, "Setting up Security for Management and Monitoring Services," on page 1183 | This section is added to address security issues. |

R.3 December 9, 2005

Page design is reformatted to comply with revised Novell documentation standards.