

User Guide

Novell® PlateSpin® Recon

3.6

March 20, 2009

www.novell.com



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. See the [Novell International Trade Services Web page \(http://www.novell.com/info/exports/\)](http://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 © 2009 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 on the [Novell Legal Patents Web page \(http://www.novell.com/company/legal/patents/\)](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 latest online documentation for this and other Novell products, see the [Novell Documentation Web page \(http://www.novell.com/documentation\)](http://www.novell.com/documentation).

Novell Trademarks

For Novell trademarks, see [the Novell Trademark and Service Mark list \(http://www.novell.com/company/legal/trademarks/tmlist.html\)](http://www.novell.com/company/legal/trademarks/tmlist.html).

Third-Party Materials

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

Contents

About This Guide	9
1 Introduction	11
1.1 About PlateSpin Recon	11
1.2 PlateSpin Recon Features	11
1.2.1 Recon Features	11
1.2.2 Planning Module Features	12
1.3 Supported Platforms	13
1.3.1 Supported Inventory and Data Collection	13
1.3.2 Supported PlateSpin Portability Suite Consolidation Sources and Targets	16
1.4 Application Scenario	17
1.5 PlateSpin Recon Best Practices	20
1.5.1 Best Practices Planning	20
1.5.2 Best Practices Installation and Setup	21
1.5.3 Best Practices Usage	22
1.5.4 Best Practices Troubleshooting	22
2 Setting Up PlateSpin Recon Licensing	23
2.1 About PlateSpin Recon Licensing	23
2.2 Activating Your Product License	23
2.2.1 Online Activation	23
2.2.2 Offline Activation	24
2.3 Managing Licenses	24
2.3.1 License Splitting	25
3 Data Center Explorer	27
3.1 About the Data Center Explorer	27
3.2 Data Center Icons and Commands	28
3.3 Working with Filters	29
3.4 Working with Groups	29
3.4.1 Creating a Group	30
3.4.2 Copying Servers to a Group	30
3.4.3 Deleting a Server from a Group	30
3.5 Working with Sites	30
3.5.1 Adding a Default Site	31
3.5.2 Deleting a Site	31
3.5.3 Modifying a Site	32
3.6 Discovering and Inventorying Data Center Resources	32
3.7 Inventory and Discovery Methods	33
3.7.1 Machines	33
3.7.2 Virtual Center	33
3.7.3 From a CSV File	34
3.7.4 Windows Domain	36
3.7.5 IP Range	37
3.7.6 Subnet	38
3.7.7 SNMP	39
3.8 Inventory Credentials	40

3.9	Discovery Options	41
3.10	Inventory Options	42
3.11	Viewing Inventory Data and Monitoring Servers	43
3.11.1	Viewing Properties	43
3.11.2	Refreshing the Inventory	44
3.11.3	Collecting Dynamic Data by Monitoring Servers	44
3.11.4	Viewing Dynamic Data	45
3.11.5	Viewing Logs	45
3.11.6	Custom Fields	45
3.12	Working with Remote Data	46
3.13	Working with PlateSpin Recon Snapshots	46
3.13.1	Exporting a PlateSpin Recon Snapshot	47
3.13.2	Importing a PlateSpin Recon Snapshot	49
3.14	Working with Data Collected by Other PlateSpin Recon Instances	52
3.15	Working with Data Collected by Other Performance Monitoring Tools	53
3.15.1	Required Performance Counters	54
3.16	Using HP Reporter Data with PlateSpin Recon	58
3.16.1	Using Formulas to Convert Metrics to Counters	58
3.16.2	Example Formulas	59
3.17	Using Performance Data Import to Import Generic Data into PlateSpin Recon	62
3.17.1	Performance Data Import Tables Data Structure	63
3.17.2	Preparing Data for Import into PlateSpin Recon	64
4	PlateSpin Recon Server and Client Tools	67
4.1	Server and Client Options	67
4.2	Job Explorer	67
4.3	Scheduled Job Explorer	68
4.4	Log Explorer	69
4.5	Credentials Manager	69
4.6	PlateSpin Analyzer	71
4.6.1	Using PlateSpin Analyzer	71
4.6.2	Inventory and Monitoring Tests	72
4.6.3	Conversion Tests	73
5	Generating and Viewing Charts and Reports	75
5.1	Chart Viewer	75
5.1.1	Chart Controls and Commands	76
5.1.2	Expression Types	77
5.1.3	Level of Detail	78
5.1.4	Data Type Statistics	78
5.1.5	Server Instance	78
5.1.6	Date Range Presets	79
5.1.7	Forecasting	79
5.2	Report Explorer	79
5.2.1	Report Explorer Template Nodes	79
5.2.2	Commands	82
5.2.3	Analysis Reports	83
5.2.4	Working with Report Templates	83
5.3	Editing Report Templates	84
5.3.1	Report Template Option Tabs	85
5.3.2	Chargeback Report Templates	89
5.3.3	Consolidation Report Templates	89
5.3.4	Consolidation Project Samples Report Templates	90
5.3.5	Diagnostics Report Templates	90

5.3.6	Disk Report Templates	90
5.3.7	Inventory Report Templates	91
5.3.8	Reclamation Report Templates	91
5.3.9	Software Report Templates	91
5.3.10	Virtual Infrastructure Report Templates	91
5.3.11	Volume Report Templates	92
5.4	View Customization	92
5.4.1	Creating and Editing Report Views	92
5.4.2	Histogram Charts	94
5.4.3	Scatter Charts	95
5.4.4	Usage Charts	96
5.5	Generating Reports	96
5.6	Document Viewer	100
5.6.1	Commands	100
5.6.2	Customizing Reports	101
5.6.3	Viewing Reports	104
5.6.4	Printing and Exporting Reports	105
5.6.5	Copying Servers from a Report to a Group	106
6	Server Consolidation	107
6.1	About Server Consolidation	107
6.2	Project Explorer	107
6.2.1	Creating Projects	109
6.2.2	Server Templates	114
6.3	Consolidation Scenarios	120
6.3.1	Determining the Minimum Number of Servers	121
6.3.2	Specifying the Number of Servers	126
6.3.3	Results of Consolidation Scenarios	129
6.3.4	Editing Scenarios	130
6.3.5	Identifying Consolidation Candidates	131
6.4	Generating Inventory Reports	131
6.4.1	Finding Servers That Meet the Consolidation Criteria	131
6.4.2	Selecting Candidates for Consolidation	131
6.5	Viewing Consolidation Reports	132
6.5.1	Project Report	133
6.5.2	Scenario Summary Report	133
6.5.3	Workload Assignment Report	135
6.5.4	Unconsolidated Workload Report	136
6.5.5	Unprotected Workload Report	136
6.6	Implementing Scenarios with Portability Suite	136
6.6.1	Implementing Phantom or Existing Target Servers	138
7	IT Chargeback	141
7.1	Understanding Chargeback	141
7.2	Using Raters	141
7.2.1	Rater Resources	142
7.2.2	Rater Charges	142
7.2.3	Virtual Server Rater Methods	143
7.3	Managing Raters	143
7.3.1	Creating a Rater	144
7.3.2	Deleting a Rater	145
7.3.3	Editing a Rater	145
7.3.4	Attaching Raters to Machines	145
7.3.5	Detaching a Rater from a Machine	146
7.4	Chargeback Reporting	146

7.4.1	How Calculated Cost is Derived	147
8	Virtual Infrastructure Management	149
8.1	Virtual Center Inventory and Monitoring	149
8.2	Scheduling and Recurrence	150
8.3	VM Sprawl Reporting	152
8.4	Virtual Machine State Reporting	152
A	Technical Reference	153
A.1	Features	153
A.2	How PlateSpin Recon Collects Data	153
A.2.1	Discovery	153
A.2.2	Inventory	154
A.2.3	Monitoring	155
A.3	Data Collection	155
A.4	Charts and Formulas	156
A.4.1	Chart Data	157
A.4.2	Utilization	158
A.4.3	Workload	158
A.5	Calculations	160
A.5.1	24-Hour Profile	160
A.5.2	Dual-Core and Multi-Core Machines	161
A.5.3	Hyper-Threading	161
A.5.4	Effective Speed	162
A.5.5	Power and Cooling	162
A.6	CPU Normalization	162
A.6.1	SPEC Data	163
A.6.2	Normalization Speed	163
A.7	PlateSpin Recon VM CPU Inventory	163
A.8	Linux/Solaris Inventory and Monitoring Commands	164
A.8.1	Inventorying Solaris	164
A.8.2	Inventorying LINUX	165
A.8.3	Monitoring	166
A.8.4	Text Processing	166

About This Guide

This documentation provides information about PlateSpin® Recon.

- ♦ Chapter 1, “Introduction,” on page 11
- ♦ Chapter 2, “Setting Up PlateSpin Recon Licensing,” on page 23
- ♦ Chapter 3, “Data Center Explorer,” on page 27
- ♦ Chapter 4, “PlateSpin Recon Server and Client Tools,” on page 67
- ♦ Chapter 5, “Generating and Viewing Charts and Reports,” on page 75
- ♦ Chapter 6, “Server Consolidation,” on page 107
- ♦ Chapter 7, “IT Chargeback,” on page 141
- ♦ Chapter 8, “Virtual Infrastructure Management,” on page 149
- ♦ Appendix A, “Technical Reference,” on page 153

Audience

This text is intended for IT staff, such as data center administrators and operators, who use PlateSpin Recon in the management and planning of the data center.

Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please take a moment to provide your comments at the [Novell Documentation Web site](http://www.novell.com/documentation/feedback) (<http://www.novell.com/documentation/feedback>).

Additional Documentation

In addition to this guide, PlateSpin Recon comes with the following documentation:

- ♦ *PlateSpin Recon Getting Started Guide*
- ♦ An online version of the *PlateSpin Recon User Guide* integrated with the product
- ♦ *Release Notes*

For help with third-party vendor products or documentation, please contact your sales or support representative for those products.

Technical Support

Contact Technical Support by:

- ♦ Telephone (North America): +1 877 528 3774 (1 877 PlateSpin). This is a toll-free number.
- ♦ Telephone (global): +1 416 203 4799. This is not a toll-free number.
- ♦ E-mail: support@platespin.com

You can also visit the [PlateSpin Technical Support Web site](http://www.platespin.com/support/) (<http://www.platespin.com/support/>)

User Forum

A user forum is available at <http://forum.platespin.com> (<http://forum.platespin.com>).

Documentation Conventions

In PlateSpin 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.

Introduction

1

This section provides information to introduce you to PlateSpin® Recon.

- ♦ [Section 1.1, “About PlateSpin Recon,” on page 11](#)
- ♦ [Section 1.2, “PlateSpin Recon Features,” on page 11](#)
- ♦ [Section 1.3, “Supported Platforms,” on page 13](#)
- ♦ [Section 1.4, “Application Scenario,” on page 17](#)
- ♦ [Section 1.5, “PlateSpin Recon Best Practices,” on page 20](#)

1.1 About PlateSpin Recon

PlateSpin Recon provides intelligence, visual analysis, and forecasting for consolidating and optimizing the data center. It collects hardware, software, and service inventory data for server workloads. PlateSpin Recon can remotely gather workload utilization statistics for a clear and concise picture of the application services running in the data center, and to determine how resources are being used.

The PlateSpin Recon Consolidation Planning Module determines the optimal fit between server resource supply and workload demand. It helps you visualize the characteristics of server workloads before and after consolidation, determine the servers required for consolidation, and expedite projects by automatically generating consolidation plans. You can also protect servers by allocating them to a suitable virtual recovery environment in case of a service outage.

PlateSpin Recon includes other features such as Virtual Infrastructure Management capabilities, as well as a resource utilization-based IT chargeback system. These can be used to help cope with the challenges presented by the virtualization and consolidation of the data center.

1.2 PlateSpin Recon Features

- ♦ [Section 1.2.1, “Recon Features,” on page 11](#)
- ♦ [Section 1.2.2, “Planning Module Features,” on page 12](#)

1.2.1 Recon Features

Remote Data Collection: PlateSpin Recon does not require the installation of agent software, eliminating the need to physically touch data center servers. All performance and inventory data is collected via standard OS instrumentation capabilities.

Rich Data Modeling: Helps you make consolidation choices based on sophisticated analysis of resources, workloads, and utilization trends. Tight integration with VMware* VirtualCenter provides greater visibility into virtual infrastructures, thereby improving data center management and operations.

Custom Report Creation and Delivery: Helps you define resource and workload parameters and generate custom visual reports for data center assessments and server consolidations. You can quickly identify consolidation candidates based on resource utilization trends, and compare workload characteristics before and after consolidation. Scheduled report delivery via e-mail or FTP ensures easy access to remote data and provides up-to-date information.

Flexible Data Capture and Export: You can export PlateSpin Recon data to a number of formats, including HTML, PDF, Word, CSV, Excel, or images. Raw data can be extracted directly from the database and delivered to business intelligence applications for advanced statistical analysis.

Enterprise-Level Scalability: Each instance of PlateSpin Recon can monitor up to 1,500 servers to ensure enterprise scalability for the world's largest data centers. Data can be aggregated from multiple PlateSpin Recon data collectors for centralized data warehousing, analysis, and planning, or to accommodate larger enterprise implementations.

Multiple-Data Center Support: You can use PlateSpin Recon in different geographical locations to remotely collect data and provide scheduled updates to a master PlateSpin Recon installation. This allows for centralized analysis and planning for initiatives such as server or data center consolidation.

Flexible Chargeback Reporting: Because virtualization creates a shared pool of computing resources, it can be difficult to manage and monitor virtual resource usage. PlateSpin Recon enables organizations to effectively allocate and share virtual resources across business units and departmental owners. Flexible chargeback reporting capabilities improve virtual infrastructure management and financial accounting by allowing organizations to accurately calculate IT costs based on actual resource usage.

VM Growth Reporting: You can generate reports to monitor the proliferation of virtual machines and avoid the administrative headaches associated with virtual infrastructure sprawl.

1.2.2 Planning Module Features

Planning: You can automatically generate server consolidation and disaster recovery plans based on detailed workload analysis to ensure the optimal fit between server workloads and virtual resources. Use forecasted data to plan for and accommodate future growth.

Workload Analysis: The PlateSpin Recon Planning Module automatically analyzes five critical workload dimensions—processor, memory, disk space, disk I/O, and network I/O—simultaneously across thousands of servers, providing consolidation plans that maximize utilization while minimizing resource contention.

Scenario Modeling: You can use custom scenarios with user-defined target server specifications, including server templates or existing virtual machine servers to create an optimal consolidation plan.

Power and Cooling Analysis: You can compare and contrast potential power and cooling cost savings and ROI derived from different consolidation scenarios. Custom fields (maintained in a central database) let you input power and cooling requirements for major hardware platforms. This feature lets organizations analyze and cost-justify green computing initiatives.

Time-Based Analysis: You can stagger multiple workloads evenly across virtual hosts and account for hourly peaks and valleys inherent in server utilization trends.

Workload and Utilization Forecasting: You can predict future workloads and resource utilization based on historical trends to better plan for server consolidation and infrastructure growth, as well as enable more proactive systems management. Forecasting data on CPU, disk, memory, and usage trends is presented in easy-to-read charts, reports, and plans.

1.3 Supported Platforms

- ♦ [Section 1.3.1, “Supported Inventory and Data Collection,” on page 13](#)
- ♦ [Section 1.3.2, “Supported PlateSpin Portability Suite Consolidation Sources and Targets,” on page 16](#)

NOTE: For information on PlateSpin Recon Client/Server supported platforms, see “*System Requirements*” in the *PlateSpin Recon Getting Started Guide*.

1.3.1 Supported Inventory and Data Collection

PlateSpin Recon supports three operating system families, Microsoft* Windows*, UNIX*/Linux* and Novell® Netware®.

- ♦ [“Supported Microsoft Windows Platforms” on page 13](#)
- ♦ [“Supported UNIX/Linux Platforms” on page 14](#)
- ♦ [“Supported Novell NetWare Platforms” on page 15](#)

Supported Microsoft Windows Platforms

For the Windows platforms, PlateSpin Recon supports the following features:

- ♦ Discovery
- ♦ Inventory
- ♦ Data Collection (Monitoring)

Table 1-1 *Supported Windows Platforms*

Version	Bit Spec
Windows Server 2000 (SP0, SP1, SP2, SP3, SP4)	32-bit
Windows Server 2003 (SP0, SP1, SP2)	32-bit, 64-bit
Windows Server 2008 (SP0)	32-bit, 64-bit
Windows Vista* (SP0, SP1)	32-bit, 64-bit
Windows NT* 4.0 Server (SP4, SP6)	32-bit
Windows NT 4.0 Server (SP6a - WMI Core version 1.5 required)	32-bit
Windows XP* Professional* (SP0, SP1, SP2, SP3)	32-bit, 64-bit
Virtual Server 2005* (version 1.0) All editions.	
Virtual Server 2005 R2 (SP1)	

Version	Bit Spec
VMware* Virtual Center* 2.0.1, 2.0.2, 2.5	
Microsoft Hyper-V	64-bit

PlateSpin Recon supports the following localized versions of Windows:

- ♦ French
- ♦ German
- ♦ Japanese

Supported UNIX/Linux Platforms

For the UNIX/Linux platforms, PlateSpin Recon supports the following features:

- ♦ Discovery
- ♦ Inventory
- ♦ Data Collection (Monitoring)

Table 1-2 *Supported UNIX/Linux Platforms*

Version	Bit Spec
Fedora* Core 3, 4, 5, 6, 7, 8, 9, 10	32-bit, 64-bit
Novell SUSE® Linux Enterprise Server (SLES) 8.0, 9.x, 10.0, 10.1, 11	32-bit, 64-bit
Xen on Novell SLES 11	32-bit, 64-bit
Novell SUSE Linux 8.0	32-bit
Novell OpenSUSE 10.1, 10.2, 10.3, 11, 11.1	32-bit, 64-bit
Citrix* XenServer* 4.1, 5.0	32-bit, 64-bit
RedHat* Linux 7.3, 8.0, 9.0, 9.1, 9.2	32-bit, 64-bit
RedHat Enterprise Linux 2.x AS/ES	32-bit
RedHat Enterprise Linux 3.0 AS/ES, 4.0 AS/ES, 5.0 AS/ES, 5.1 AS/ES	32-bit, 64-bit
Ubuntu* 7.1, 8.04, 8.10	32-bit, 64-bit
CentOS* 2.0	32-bit
CentOS 3.9, 4.6, 5.1, 5.2	32-bit, 64-bit
Sun* Solaris* 8, 9, 10 (SPARC*/x86)	32-bit, 64-bit
Sun Solaris 7 (SPARC)	32-bit
Sun Solaris 10 Zones on Sun Solaris 10 Container	32-bit, 64-bit
VMware ESX* Server 2.5.3, 2.5.4, 2.5.5, 3.0.0, 3.0.1, 3.0.2, 3.0.3, 3.5.0, 3i	32-bit, 64-bit
VMware GSX* Server 3.2 (for Windows)	32-bit, 64-bit
VMware 1.x+ (for Windows)	32-bit, 64-bit

Supported Novell NetWare Platforms

For the Novell NetWare platforms, PlateSpin Recon supports the following features:

- ♦ Discovery
- ♦ Inventory
- ♦ Data Collection (Monitoring)

Table 1-3 *Supported Novell NetWare Platforms*

Version	Bit Spec
Novell NetWare 5.1 (SP8)	32-bit, 64-bit
Novell NetWare 6.0 (SP5)	32-bit, 64-bit
Novell NetWare 6.5 (SP6)	32-bit, 64-bit

IMPORTANT: and Download and install the following Novell Clients on the PlateSpin Recon server computer to enable monitoring of NetWare computers. The client downloads are free. When installing, select *Custom Installation* and install the Novell Modular Authentication Service (NMAS) client.

- ♦ Novell Client 4.91 SP5 (<http://download.novell.com/Download?buildid=qmMAWSRy5q4~>)
- ♦ Novell Client 4.91 Post-SP5 NWFS.SYS (<http://download.novell.com/Download?buildid=m7NbjoMFSxU&donotredirect=true>)

Installation of the Novell Client impacts the standard Windows logon. See the Novell Client documentation for more information.

Only the CPU and disk utilization counters are supported for Netware 5.1 and 6.0. Only the CPU, Disk and memory utilization counters are supported for Netware 6.5 and higher.

1.3.2 Supported PlateSpin Portability Suite Consolidation Sources and Targets

Consolidation scenarios for workloads with the following operating systems and targets can be implemented by using PlateSpin Portability Suite:

Table 1-4 *Portability Suite Supported Source/Target Operating Systems*

Source Operating System	Target Operating System
♦ Microsoft Windows Server 2003	♦ VMware Virtual Infrastructure 3.0 - 3.0.1
♦ Microsoft Windows Server 2000	♦ VMware ESX Server 2.5.x
♦ Microsoft Windows Server 2000 Advanced	♦ VMware Server 1.0, 1.0.1 - 1.0.3 for Windows
♦ Microsoft Windows XP	♦ VMware GSX Server 3.1, 3.2 for Windows
♦ Microsoft Windows NT	♦ Microsoft Virtual Server 2005
♦ Red Hat Linux 7.3, 8.0	♦ XenEnterprise*
♦ Red Hat Enterprise Linux AS 3.0, ES 3.0	♦ Virtual Iron*
♦ Red Hat Enterprise Linux AS 4.0, ES 4.0	♦ Citrix XenServer
♦ Novell SLES 9.0	♦ Sun Solaris 10 Zones
	♦ Xen on Novell SLES 11

For more detailed information about source/target operating systems for PlateSpin Portability Suite conversions, see Knowledge Base article [Q20002 \(http://support.platespin.com/kb2/article.aspx?id=20002\)](http://support.platespin.com/kb2/article.aspx?id=20002).

1.4 Application Scenario

PlateSpin Recon and PlateSpin Portability Suite combine to form the only suite that automates the entire capacity planning and server consolidation project from start to finish. The following table illustrates how to assess, design, and implement a solution for optimizing and consolidating the data center.

Table 1-5 *Data Center Optimization and Consolidation Workflow*

Action	Process	Description	Solution
Assess	1. Inventory	<ul style="list-style-type: none">♦ Remotely discover server hardware and software assets across the data center♦ Gather detailed information for each server such as operating system, installed and running applications or services, patch levels, CPU, memory, network and disk resources♦ Connect to the network, then select relevant computers or domains; data on hardware and resources is collected in a centralized database♦ Organize servers by predefined and user-defined groupings for organized analysis	PlateSpin Recon: Remotely and agentlessly identify server resources and workloads in the data center and create a complete picture of hardware and software inventory, utilization levels, and workload sizes

Action	Process	Description	Solution
Assess	2. Collect	<ul style="list-style-type: none"> ♦ Identify workloads and collect utilization data over days, weeks or months to determine workload sizes and appropriate resource matching ♦ Create dynamic reports to identify workload and resource mismatches and identify candidate workloads for consolidation or resource upgrading ♦ Rank and sort server workloads by CPU, network, disk, memory inventory, and/or performance ♦ View summary workload sizing data or drill down to point-in-time granular levels ♦ Collect and analyze data from multiple sites to create a more complete picture of the data center ♦ Export the data to third-party tools or copy visual graphs to business presentations 	PlateSpin Recon: Remotely and agentlessly identify server resources and workloads in the data center and create a complete picture of hardware and software inventory, utilization levels, and workload sizes

Action	Process	Description	Solution
Design	3. Analyze	<ul style="list-style-type: none"> ◆ Identify virtualization candidates ◆ Create scenarios for distributing workloads across servers to maximize utilization ◆ Balance application workloads over time and across hardware resources to minimize resource contention ◆ Use “what-if” modeling to determine different combinations of hardware and virtual hosts required for the project ◆ Generate consolidation or disaster recovery plans with detailed project, scenario, and workload assignment reports and charts ◆ Compare different scenarios based on total cost of ownership, consolidation ratio, rack space, and power needs 	PlateSpin Recon: Generate consolidation scenario plans to optimally combine and allocate workloads to new virtual hosts
Implement	4. Consolidate	<ul style="list-style-type: none"> ◆ Accelerate the server consolidation project by streaming physical servers into virtual environments such as VMware ESX Server, VMware Server, and Microsoft Virtual Server ◆ Automate migrations between dissimilar physical servers or from physical servers to blades 	PlateSpin Portability Suite: Automatically stream data, applications, and operating systems to any physical or virtual platform over the network

Action	Process	Description	Solution
Manage	5. Monitor	<ul style="list-style-type: none"> ♦ Keep track of workload performance ♦ Identify over- and under-utilized workloads for ongoing optimization 	PlateSpin Recon
	6. Chargeback	<ul style="list-style-type: none"> ♦ Assign customized rates to machines ♦ Measure resource usage and bill business owners for usage ♦ Create reports on usage and cost 	
	7. Manage Virtual Infrastructure	<ul style="list-style-type: none"> ♦ Monitor VM creation ♦ Report on VM growth 	

1.5 PlateSpin Recon Best Practices

These Best Practices are designed for data center operators and administrators who use PlateSpin Recon to monitor workloads and do consolidation planning based on collected data. The information provided here has been compiled by the PlateSpin support team based on experience with hundreds of Real-Time Consolidation Accelerator users.

Where applicable, guidelines in Best Practices refer to PlateSpin Support Knowledge Base articles. To access the articles, use your PlateSpin download credentials. If you have forgotten your password, you can request a password reminder email at: <http://www.platespin.com/downloads/forgotpassword.aspx> (<http://www.platespin.com/downloads/forgotpassword.aspx>)

- ♦ [Section 1.5.1, “Best Practices Planning,” on page 20](#)
- ♦ [Section 1.5.2, “Best Practices Installation and Setup,” on page 21](#)
- ♦ [Section 1.5.3, “Best Practices Usage,” on page 22](#)
- ♦ [Section 1.5.4, “Best Practices Troubleshooting,” on page 22](#)

1.5.1 Best Practices Planning

Guideline	Knowledge Base Article
Ensure that servers you plan to monitor are running on platforms and operating systems supported by PlateSpin Recon, and that they meet all software requirements.	<p>Q20909 (http://support.platespin.com/kb2/article.aspx?id=20909): Supported platforms for Inventory and Data Collection</p> <p>Q20729 (http://support.platespin.com/kb2/article.aspx?id=20729): What PlateSpin Recon requires to add a server to its inventory</p>

Guideline	Knowledge Base Article
Plan your project and define the following parameters:	
<ul style="list-style-type: none"> ♦ Number of servers to monitor ♦ Duration of monitoring period 	
Order the proper PlateSpin Recon licenses based on project parameters defined above.	Q20876 (http://support.platespin.com/kb2/article.aspx?id=20876) : PlateSpin License Entitlement Manager
Use the License Entitlement Manager Web portal to view license orders and manage activation codes.	

1.5.2 Best Practices Installation and Setup

Guideline	Knowledge Base Article
Ensure that the designated PlateSpin Recon Server is compliant with the minimum system requirements.	Q20904 (http://support.platespin.com/kb2/article.aspx?id=20904) : System Requirements for PlateSpin Recon
Consider recommended hardware specifications for various project sizes. PlateSpin Recon 3.5 supports up to 2,000 workloads per Server.	Q20887 (http://support.platespin.com/kb2/article.aspx?id=20887) : PlateSpin Recon Installation Best Practices Guide
For large-scale projects (2,000 - 10,000 workloads) use multiple PlateSpin Recon collectors. Centralize analysis and planning by automatically and securely synchronizing data from collectors to a secure master instance of PlateSpin Recon.	
Ensure that the PlateSpin Recon Server has enough disk space for the projected database growth based on the number of workloads you plan to monitor.	Q20910 (http://support.platespin.com/kb2/article.aspx?id=20910) : How large the PlateSpin Recon Database will grow
Set up designated PlateSpin Recon user accounts with administrative rights.	Q20726 (http://support.platespin.com/kb2/article.aspx?id=20726) : PlateSpin Recon Best Practices for User Accounts
When using PostgreSQL*, set it up to allow remote connections (not allowed by default).	Q20889 (http://support.platespin.com/kb2/article.aspx?id=20889) : Configure PostgreSQL to allow Remote connections
To inventory virtual machines hosted by VMware GSX Server, install the VmCOM API on the PlateSpin Recon server (otherwise the VMs will not be visible).	Q20516 (http://support.platespin.com/kb2/article.aspx?id=20516) : Installing the VMware GSX Server VmCOM API on the PlateSpin Recon server.

1.5.3 Best Practices Usage

Guideline	Knowledge Base Article
Use this generic consolidation planning workflow:	
Step 1: Inventory the servers.	
Step 2: Monitor the servers and collect data.	
Step 3: Analyze the collected data and create consolidation projects and scenarios by using PlateSpin Recon charts and the Consolidation Planning Module.	
When you have finished, consolidate the servers by streaming physical servers into virtual environments by using PlateSpin PowerConvert.	
For early detection of potential inventory and monitoring issues, use PlateSpin Recon PlateSpin Analyzer to scan your network before starting the inventory.	Q20825 (http://support.platespin.com/kb2/article.aspx?id=20825) : Using PlateSpin Analyzer and understanding its results.
If possible, use Microsoft Operations Manager (MOM) 2005 and set up PlateSpin Recon to connect and collect monitoring data directly from it. This feature allows you to reduce your project timeline from weeks to days.	Q20903 (http://support.platespin.com/kb2/article.aspx?id=20903) : PlateSpin Recon with MOM 2005

1.5.4 Best Practices Troubleshooting

Problem	Knowledge Base Article
Issues when attempting to add servers to the PlateSpin Recon Inventory, including the following error messages: Network Path Not found Access Denied The RPC Server is unavailable Failed. The Network location cannot be reached	Q20525 (http://support.platespin.com/kb2/article.aspx?id=20525) : Troubleshooting problems when adding servers to the PlateSpin Recon Inventory
When using MSDE 2000, PlateSpin Recon stops logging data after a certain period of time.	Q20556 (http://support.platespin.com/kb2/article.aspx?id=20556) : PlateSpin Recon stopped logging data
When attempting to import a machine from a PlateSpin Recon snapshot, the following error is displayed: Cannot import MachineObject 'MACHINE-NAME [MACHINE-NAME]'	Q20900 (http://support.platespin.com/kb2/article.aspx?id=20900) : Cannot import MachineObject error when importing a machine from a snapshot

Setting Up PlateSpin Recon Licensing

2

This section provides information about PlateSpin® Recon licensing.

- ♦ [Section 2.1, “About PlateSpin Recon Licensing,” on page 23](#)
- ♦ [Section 2.2, “Activating Your Product License,” on page 23](#)
- ♦ [Section 2.3, “Managing Licenses,” on page 24](#)

2.1 About PlateSpin Recon Licensing

A PlateSpin Recon license entitles the user to one instance of PlateSpin Recon Server. No license is required to install and use the PlateSpin Recon Client.

PlateSpin Recon provides several licensing options at the time of purchase.

Per Use: Limits the amount of time and the number of servers that can be monitored. For example, ten server-days allows monitoring one server for ten days or two servers for five days.

Per Core: Licenses are tied to physical cores for monitoring, regardless of the operating system, platform or the number of servers involved. When available, virtual machines consume licenses based on the cores of their host. If the host has not been inventoried, licenses are consumed based on the number of virtual CPUs assigned to the VM.

NOTE: License types cannot be combined. Licensing does not apply limits during discovery, but only during inventory and monitoring.

For more information about Virtual Center licensing, see [Section 3.7.2, “Virtual Center,” on page 33](#).

2.2 Activating Your Product License

After PlateSpin Recon is purchased, PlateSpin sends an e-mail containing an activation code and product details. After installation, launch PlateSpin Recon and click *Activate* to launch the License Wizard. It prompts you to choose either online or offline activation.

- ♦ [Section 2.2.1, “Online Activation,” on page 23](#)
- ♦ [Section 2.2.2, “Offline Activation,” on page 24](#)

2.2.1 Online Activation

Online activation requires Internet access. An activation code and the e-mail address that was used to download the product are also necessary.

NOTE: HTTP proxy might cause failures during online activation. Offline activation is recommended for users in HTTP proxy environments.


To activate the license online:

- 1 Select the *Online Activation* option and click *Next*.
- 2 Specify the e-mail address and activation code.
- 3 Click *Next* to obtain the PlateSpin Recon license.

The PlateSpin Recon Client obtains the required license via the Internet and activates the PlateSpin Recon Server.

2.2.2 Offline Activation

A user name, password, activation code, e-mail address, and hardware ID are required to get the license key file from the [PlateSpin Activation site \(http://www.platespin.com/productactivation/ActivateOrder.aspx\)](http://www.platespin.com/productactivation/ActivateOrder.aspx).



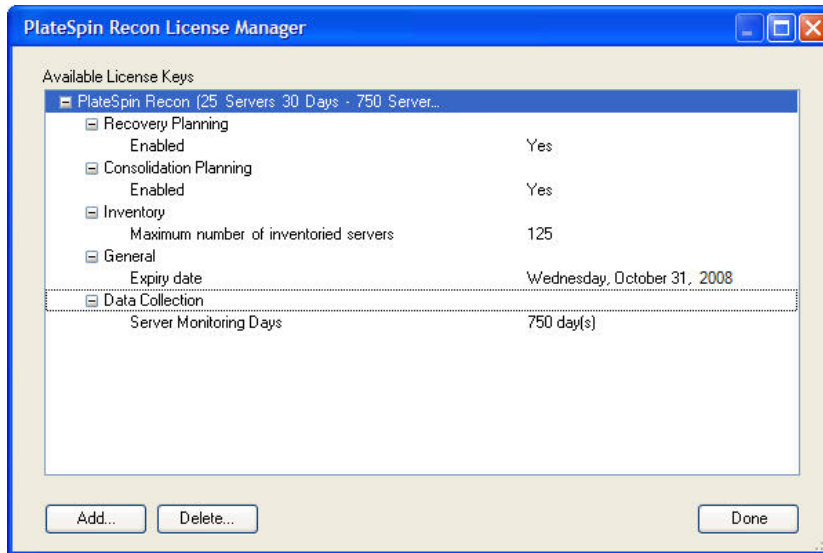
To activate the license offline:

- 1 Select the *Offline Activation* option and click *Next*.
- 2 When prompted, enter the PlateSpin Recon license key file location.

The PlateSpin Recon Client activates the PlateSpin Recon Server based on the license key file.

2.3 Managing Licenses

To access PlateSpin Recon License Manager, navigate to *Tools* and click *License Manager* on the PlateSpin Recon Client tool bar. All PlateSpin Recon licenses are displayed in the PlateSpin Recon License Manager window.



Additional licenses can be added at any time by clicking *Add*. Expired licenses can be deleted by selecting them and clicking *Delete*.

2.3.1 License Splitting

There are cases when more than one PlateSpin Recon Server is needed, such as a system using a number of remote sites that provide distributed monitoring to a central site for analysis, or a very large data center that exceeds the capacity of one PlateSpin Recon Server.

If necessary, a purchased license can be split between 2 or more servers. For example, a per-server license to monitor 1000 servers can be split into one license to monitor 400 servers and another to monitor 600. A per-use license for 3000 server-days of monitoring can be split into one license for 1200 server-days and one license for 1800 server-days.

For more information about splitting licenses, refer to Knowledge Base article [Q20876 \(http://support.platespin.com/kb2/article.aspx?id=20876\)](http://support.platespin.com/kb2/article.aspx?id=20876).

Data Center Explorer

3

This section provides information about the PlateSpin® Recon Data Center Explorer and the features accessible from it.

- ♦ [Section 3.1, “About the Data Center Explorer,” on page 27](#)
- ♦ [Section 3.2, “Data Center Icons and Commands,” on page 28](#)
- ♦ [Section 3.3, “Working with Filters,” on page 29](#)
- ♦ [Section 3.4, “Working with Groups,” on page 29](#)
- ♦ [Section 3.5, “Working with Sites,” on page 30](#)
- ♦ [Section 3.6, “Discovering and Inventorying Data Center Resources,” on page 32](#)
- ♦ [Section 3.12, “Working with Remote Data,” on page 46](#)

3.1 About the Data Center Explorer

The Data Center Explorer is one of several major PlateSpin Recon *views*. It is used to organize, monitor, view, and manage information about data center assets. Data center assets are physical and virtual machines, and physical and virtual workloads. PlateSpin Recon uses information about them to display dynamic charts, generate performance reports, and generate performance forecasts.

The Data Center Explorer is the central point for:

- ♦ Interacting with physical and virtual machines in the data center
- ♦ Discovering and inventorying hardware and workloads
- ♦ Adding, monitoring, filtering, grouping, or deleting servers

In the Data Center Explorer, physical and virtual servers are represented in a hierarchical tree structure.

Table 3-1 *Data Center Explorer Nodes*

Nodes	Description
All	Lists all systems
Discovered	Lists servers that have been discovered but not yet inventoried
Filters	Contains servers listed by built-in criteria, such as domain, operating system, or monitoring status
Groups	Contains manually grouped servers according to any specific data segregation requirements
Sites	Contains servers by data site, including a default local site and others you create to work with data about servers monitored by other instances of PlateSpin Recon, or imported from third-party performance monitoring tools

3.2 Data Center Icons and Commands

Table 3-2 *Data Center Explorer Icons*











	Physical Server		Virtual Center Folder
	Virtual Machine		Virtual Center Cluster
	Groups		Virtual Center Pool
	Virtual Center		Primary Site
	Virtual Center Data Center		Secondary Site

Table 3-3 *Data Center Explorer Icon Status Indicators*





















	Monitored		Virtual Machine Powered On
	Monitored virtual machine		Virtual Machine Powered Off
	Not monitored		Unresponsive/error
	Not inventoried virtual machine		Warning

Table 3-4 *Data Center Explorer Commands and Associated Icons*

Commands	Submenu	Description
 Inventory		Inventories the selected servers . The selection must be discovered but not yet inventoried. To inventory servers that have already been inventoried, use <i>Refresh Inventory</i> .
 Start Monitoring		Monitors the selected servers . The selection can be a node (<i>All, Filters, Groups, or Sites</i>) in the Data Center Explorer pane, or a particular Filter, Group, or Site to monitor all child servers.
 Stop Monitoring		Stops monitoring the selected servers . The selection can be a node (<i>All, Filters, Groups, or Sites</i>) in the Data Center Explorer pane, or a particular Filter, Group, or Site to stop monitoring all child servers.
 Logs	View Logs	View logs of selected servers .
	Acknowledge Logs	If there are some errors associated with a machine, an error icon is displayed on that machine. If the errors are acceptable, you can use <i>Acknowledge Logs</i> to remove the error icon from the display of the machine.
 Create Group		Create a new group of servers.

Commands	Submenu	Description
 Copy to Group	<groupname>	Copies the selected servers to the designated group.
 Rename Group		Renames the selected group.
 Create Template		Opens the New Server Template dialog box.
 Attach Credentials		Opens the Attach Credentials dialog box.
Launch	Terminal Services	Opens the selected application (if it is installed). Paths are configurable by selecting <i>Options</i> from the <i>Tools</i> menu, then clicking the <i>Client</i> tab.
	SSH	
	VMware Remote Console	
	VMware Virtual Infrastructure Client	
	VMware Web Access	
 Delete		Delete the server or group of servers.
Chargeback		Opens the Raters for <i>name</i> dialog box.
View XML		Displays the selected machine details as an XML document. Only the first machine is displayed if there are multiple selections.
 Properties		Opens a Properties dialog box for the selected server.
 Set Custom Fields		Opens a Properties dialog box (Custom Fields page) for the selected servers when multiple servers are selected.

To apply Data Center Explorer commands to servers, right-click a server and select the required command from the menu.

3.3 Working with Filters

The *Filters* node contains additional nodes that list discovered and inventoried servers organized into logical groups based on domain, server make and model, operating system, status, subnet mask, virtual machine servers, and workgroup. Filters cannot be added or removed, and servers cannot be added to filters.

3.4 Working with Groups

The *Groups* node contains user-defined collections of servers. Use groups to logically group servers by location, function, application, and so on. Groups can be created, deleted, and nested. Servers can belong to more than one group.

- ♦ [Section 3.4.1, “Creating a Group,” on page 30](#)
- ♦ [Section 3.4.2, “Copying Servers to a Group,” on page 30](#)
- ♦ [Section 3.4.3, “Deleting a Server from a Group,” on page 30](#)

3.4.1 Creating a Group

- 1 In the Data Center Explorer, right-click *Groups* and click *Create Group*.

The Create Group dialog box is displayed.

- 2 Type a *Group Name*.

- 3 Click *OK*.

The group is created and ready for servers to be copied to it.

3.4.2 Copying Servers to a Group

- 1 Select one or more servers in the Data Center Explorer. Right-click the selection and click *Copy to Group*.

You can also drag servers from under the *All*, *Filtered* or *Sites* node and drop them on an existing group to copy them to that group.

- 2 Click the desired group on the submenu that is displayed.

The server is displayed under the selected group in the Data Center Explorer.

3.4.3 Deleting a Server from a Group

- 1 Right-click a server in the Data Center Explorer and click *Delete*.

A confirmation dialog box is displayed.

- 2 (Optional) Select *Delete <server> and associated monitoring data from the database* to permanently remove the monitoring data for the server when it is deleted from the group.

- 3 Click *Delete*.

You can also press the Del key.

The server is removed from the group in the Data Center Explorer.

3.5 Working with Sites

The *Sites* node contains logically organized sets of servers for analysis and reporting on data collected from multiple locations.

Each PlateSpin Recon installation includes a default *Local* site that contains machines that you work with using the local instance of PlateSpin Recon.

You can work with data external to the local PlateSpin Recon instance, in which case you must first create a new site of the appropriate site type.

NOTE: Machines imported from a remote site can be viewed in reports and charts once the data has been synchronized. Further inventorying or monitoring cannot take place, though. To do this, perform the actions at the original remote site and then import the machines again. You will also need to synchronize the data again.

Table 3-5 *Site Types*

Type	Description
Local	To work with PlateSpin Recon snapshots (*.prs files) exported from other instances of PlateSpin Recon. When it is exported from the local instance of PlateSpin Recon, a snapshot can be imported into the default Local site. See Section 3.13, “Working with PlateSpin Recon Snapshots,” on page 46.
Remote PlateSpin Recon	To connect to and work with data collected by a remote PlateSpin Recon instance. See Section 3.14, “Working with Data Collected by Other PlateSpin Recon Instances,” on page 52.
Remote	To work with data collected by supported third-party performance monitoring tools, such as Microsoft Operations Manager 2005, Windows Performance Monitor (perfmon.exe), HP Reporter, and other generic performance data gathering tools. See Section 3.15, “Working with Data Collected by Other Performance Monitoring Tools,” on page 53.

Using PlateSpin Recon sites, you can designate a PlateSpin Recon Server as the Master Server and carry out all analysis and reporting tasks within a single PlateSpin Recon environment.

- ♦ [Section 3.5.1, “Adding a Default Site,” on page 31](#)
- ♦ [Section 3.5.2, “Deleting a Site,” on page 31](#)
- ♦ [Section 3.5.3, “Modifying a Site,” on page 32](#)

3.5.1 Adding a Default Site

- 1 Right-click anywhere in the Data Center Explorer and choose *Sites > Add Site*.
The Add Site dialog box is displayed.
- 2 Type a unique *Name* for the site.
- 3 Ensure that *Local* is selected in the *Type* drop-down box.
- 4 (Optional) Specify a *Description* for the site.
- 5 Ensure that the *Run Import Site Snapshot Wizard* check box is deselected.
- 6 Click *OK*.

The site is created and is displayed under the *Sites* node in the Data Center Explorer.

For information on adding Remote or Remote PlateSpin Recon sites, see [Section 3.12, “Working with Remote Data,” on page 46.](#)

3.5.2 Deleting a Site


- 1 In the Data Center Explorer, right-click the site you want to delete and select *Delete*.
A Confirm Delete dialog box is displayed.
- 2 Click *Yes*.
PlateSpin Recon removes the selected site and associated resources from the Data Center Explorer.

3.5.3 Modifying a Site

- 1 In the Data Center Explorer, right-click the site you want to modify and select *Sites > Edit Site*. The Edit Site dialog box is displayed.
- 2 Make the required modifications.
- 3 Click *OK*.








3.6 Discovering and Inventorying Data Center Resources

You add servers to the Data Center Explorer by using the *Inventory* button



There are seven inventory methods available:

Table 3-6 *Inventory Methods*

Discover and Inventory	Discover and Optionally Inventory
 Machines (page 33)	 Windows Domain (page 36)
 Virtual Center (page 33)	 IP Range (page 37)
 From a CSV File (page 34)	 Subnet (page 38)
	 SNMP (page 39)

Each inventory method is unique and has different requirements. See [Section 3.7, “Inventory and Discovery Methods,” on page 33](#) for more information.

PlateSpin Recon gets information about data center assets in two stages:

- ♦ **Discover:** Discovering machines in the network retrieves information such as hostnames, IP addresses, MAC addresses, and port information.
- ♦ **Inventory:** Inventorying the previously discovered machines retrieves deeper data for each server, such as operating system, installed applications, services, processor speed, memory, network, and disk resources.

After machines are inventoried, they can be monitored and included in all report types, regardless of which method is used for inventory.

For information about ports used by PlateSpin Recon to inventory servers, see Knowledge Base article [Q20571 \(http://support.platespin.com/kb2/article.aspx?id=20571\)](http://support.platespin.com/kb2/article.aspx?id=20571).

A common way to use discovery and inventory is to discover a large number of servers on the network. These servers appear under the *Discovered* node in the Data Center Explorer. Next, select the discovered servers and inventory them. Finally, select which inventoried servers you want to monitor. It is also possible to discover, inventory, and start monitoring servers in a single step.

3.7 Inventory and Discovery Methods

The following are the inventory and discovery methods available when adding servers to the Data Center Explorer.

- ♦ [Section 3.7.1, “Machines,” on page 33](#)
- ♦ [Section 3.7.2, “Virtual Center,” on page 33](#)
- ♦ [Section 3.7.3, “From a CSV File,” on page 34](#)
- ♦ [Section 3.7.4, “Windows Domain,” on page 36](#)
- ♦ [Section 3.7.5, “IP Range,” on page 37](#)
- ♦ [Section 3.7.6, “Subnet,” on page 38](#)
- ♦ [Section 3.7.7, “SNMP,” on page 39](#)

3.7.1 Machines

Used to add physical and virtual servers.

To add machines, specify the *Hostname or IP of the machine(s) to inventory*. Multiple hostnames/IPs can be entered, but they need to all use the same credentials. Separate each entry with a comma.

Inventory is mandatory when adding machines, although you can use the Inventory drop-down to perform the inventory at a later, scheduled time.

To enable the *Inventory* button, you need to select credentials for the machines you’re adding. For more information, see [Section 3.8, “Inventory Credentials,” on page 40](#).

You can also start monitoring the machines when they are inventoried and/or add them to existing groups. For more information, see [Section 3.10, “Inventory Options,” on page 42](#).

3.7.2 Virtual Center

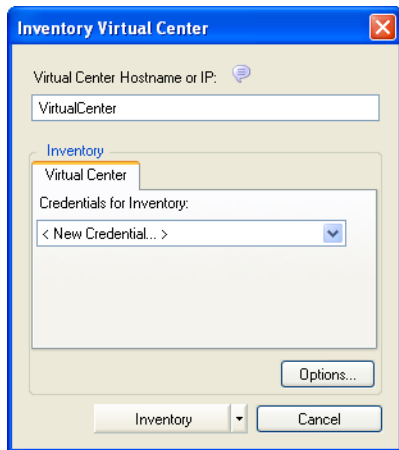
Used to add Virtual Center servers.

To add Virtual Center servers, specify the *Virtual Center Hostname or IP*. Inventorying a Virtual Center involves collecting information about all the servers and virtual machines in the Virtual Center. This is much more efficient than discovering and inventorying all the servers and virtual machines individually.

Inventory is mandatory when adding Virtual Center servers, although you can use the Inventory drop-down list to perform the inventory at a later, scheduled time.

To enable the *Inventory* button, you need to select credentials for the Virtual Center servers you’re adding.

Figure 3-1 *Inventory Virtual Center Dialog Box*



Virtual Center contents appear in the Data Center Explorer the same way they are organized in the Virtual Center, with all nodes, clusters, resource pools, and virtual machine nodes. The machines do not appear under *All* in the Data Center Explorer until they are inventoried.

When using per core licensing, you must first inventory a Virtual Center before inventorying ESX Servers that belong to it, or the virtual machines on those ESX Servers. For more information, see [“Per Core:” on page 23](#).

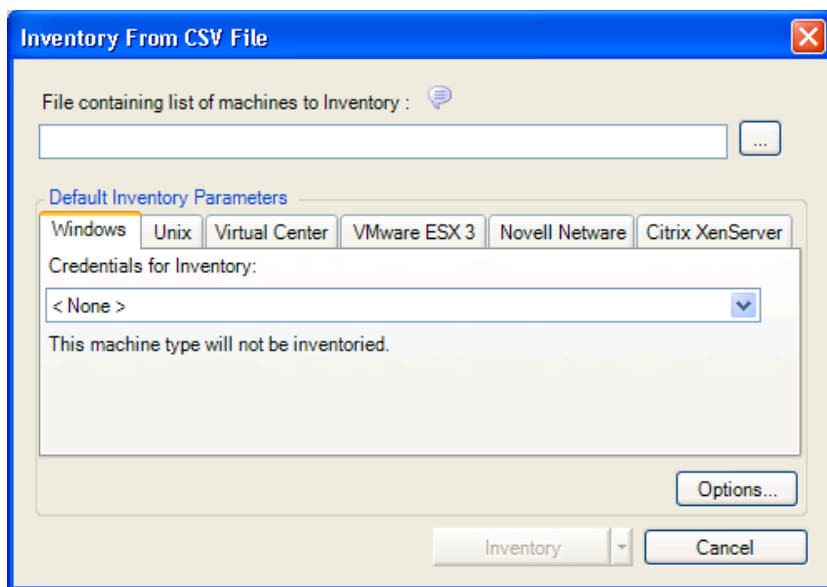
You can also start monitoring the servers when they are inventoried and/or add them to existing groups. For more information, see [Section 3.10, “Inventory Options,” on page 42](#).

3.7.3 From a CSV File

Used to add servers whose information resides in a comma-separated value file.

To inventory servers from a CSV file, type the path (or use the adjacent browse button) to locate a CSV file containing servers to inventory.

Figure 3-2 *Inventory From CSV File Dialog Box*



NOTE: The *Default Inventory Parameters* area is disabled if the CSV file contains credentials for all its machines.

Inventory is mandatory when adding servers from a CSV file, although you can use the Inventory drop-down list to perform the inventory at a later, scheduled time.

To enable the *Inventory* button, you need to select credentials for the machines you're adding. For more information, see [Section 3.8, "Inventory Credentials," on page 40](#).

Listings in the CSV file must use the following format: *Address, [Username], [Password], [Description]*.

Table 3-7 *CSV File Requirements*

Field	Required	Description
Address	Yes	IP address or hostname of the machine.
Username	[Optional]	Username for the credential. If this is left blank, the default credentials are used.
Password	[Optional]	Password of the credential. This is required only when creating a credential.
Description	[Optional]	Description for the credential.

Field	Required	Description
[inventory type]	[Optional]	<p>Can be any of the following values. Values must match case and spacing as shown below:</p> <ul style="list-style-type: none"> ♦ UNIX ♦ VirtualCenter ♦ VMwareESX3 ♦ NetWare ♦ CitrixXenServer ♦ Windows <p>If the value is unspecified, Windows and UNIX are the defaults, determined by the format of the credentials supplied:</p> <ul style="list-style-type: none"> ♦ UNIX - [user]/[password] ♦ Windows - [domain]\[user]/password <p>This field is also expected to be at the end of the CSV file.</p>

You can also start monitoring the machines when they are inventoried and/or add them to existing groups. For more information, see [Section 3.10, “Inventory Options,” on page 42](#).

3.7.4 Windows Domain

Used to add all the servers discovered in the specified Windows Domain.

Type a Windows Domain name and use the drop-down list to specify which machines to discover: *All Machines*, *Servers*, or *Workstations* (NT4 Workstation, Windows XP, Windows Vista).

Figure 3-3 Domain Discovery and Inventory Dialog Box

The screenshot shows the 'Domain Discovery and Inventory' dialog box. It is divided into two main sections: 'Discovery' and 'Inventory'.
In the 'Discovery' section:
- 'Enter the name of the Windows Domain:' is a text box containing 'MyDomain'.
- 'Machines to discover:' is a dropdown menu set to 'All'.
- 'Organizational Unit Filter:' is an empty text box.
- 'Domain Credentials' is a dropdown menu set to 'MyDomain\MyName'.
- An 'Options...' button is located at the bottom right of this section.
In the 'Inventory' section:
- The 'Inventory' checkbox is checked.
- There are tabs for 'Windows', 'Unix', 'Virtual Center', 'VMware ESX 3', 'Novell Network', and 'Citrix XenServer'.
- 'Credentials for Inventory:' is a dropdown menu set to '< None >'.
- Below this is the text 'This machine type will not be inventoried.'
- An 'Options...' button is located at the bottom right of this section.
At the very bottom of the dialog are 'Inventory' and 'Cancel' buttons.

NOTE: Clear the *Inventory* check box to only discover machines.

Discovery options can be configured by clicking the *Options* button beneath the Domain Credentials drop-down list. For more information, see [Section 3.10, “Inventory Options,” on page 42](#).

An Organizational Unit (OU) filter can be used to focus discovery on a specific part of a domain. For more information, see [“Domain Discovery:” on page 153](#).

To enable the *Inventory* button, you need to select credentials for the machines you’re adding. For more information, see [Section 3.8, “Inventory Credentials,” on page 40](#).

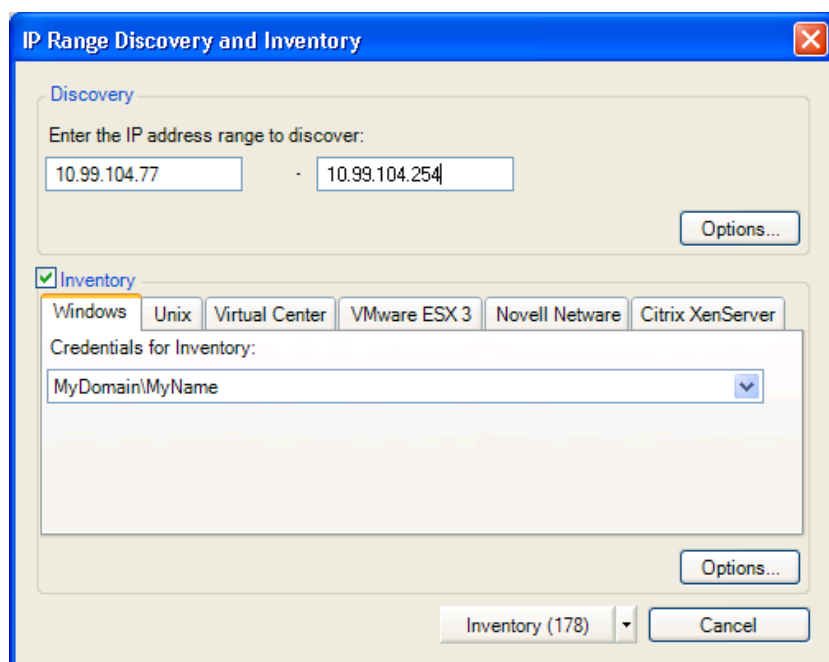
You can also start monitoring the machines when they are inventoried and/or add them to existing groups. For more information, see [Section 3.10, “Inventory Options,” on page 42](#).

3.7.5 IP Range

Used to add servers within a specified IP address range.

Enter the IP address range of the servers to discover.

Figure 3-4 IP Range Discovery and Inventory Dialog Box



NOTE: Clear the *Inventory* check box to run a discovery operation without inventorying server details.

Discovery options can be configured by clicking the *Options* button in the Discovery area. For more information, see [Section 3.10, “Inventory Options,” on page 42](#).

To enable the *Inventory* button, you need to select credentials for the machines you’re adding. For more information, see [Section 3.8, “Inventory Credentials,” on page 40](#). The numbers that are displayed in the *Inventory* button indicate how many individual IP addresses that are contained within the range you entered.

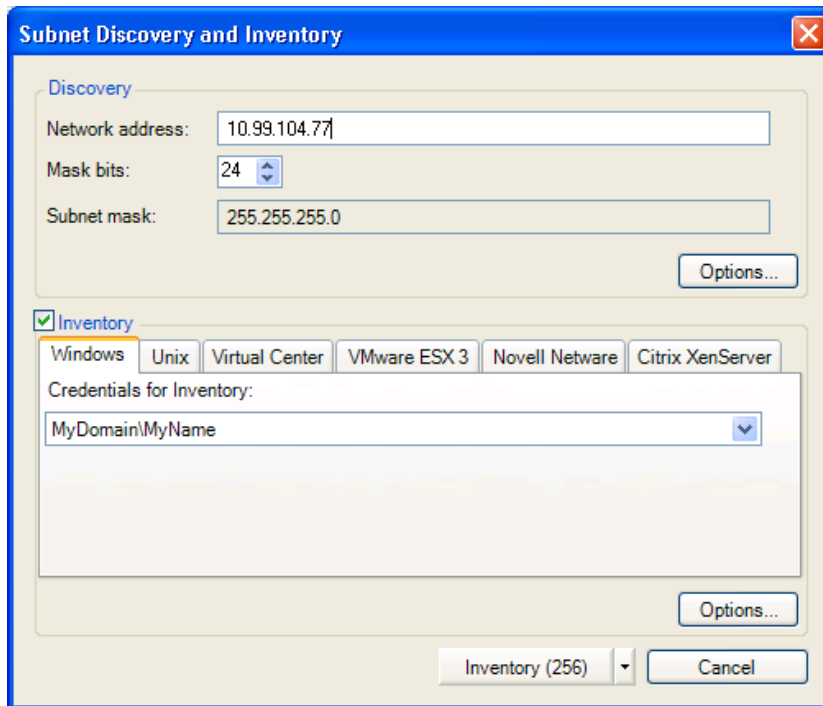
You can also start monitoring the machines when they are inventoried and/or add them to existing groups. For more information, see [Section 3.10, “Inventory Options,” on page 42](#).

3.7.6 Subnet

Used to add servers discovered in the specified subnet.

Type the *Network Address* and the *Mask bits* of the subnet to be discovered.

Figure 3-5 Subnet Discovery and Inventory Dialog Box

The image shows a Windows-style dialog box titled "Subnet Discovery and Inventory". It has a blue title bar with a close button. The dialog is divided into two main sections: "Discovery" and "Inventory". The "Discovery" section contains three text input fields: "Network address:" with the value "10.99.104.77", "Mask bits:" with a spinner set to "24", and "Subnet mask:" with the value "255.255.255.0". There is an "Options..." button to the right of these fields. The "Inventory" section is below, with a checked checkbox. It has a tabbed interface with tabs for "Windows", "Unix", "Virtual Center", "VMware ESX 3", "Novell Netware", and "Citrix XenServer". The "Windows" tab is selected. Below the tabs is a "Credentials for Inventory:" label and a dropdown menu showing "MyDomain\MyName". There is another "Options..." button to the right. At the bottom of the dialog, there is a button labeled "Inventory (256)" and a "Cancel" button.

NOTE: Clear the *Inventory* check box to run a discovery operation without inventorying server details.

Discovery options can be configured by clicking the *Options* button in the Discovery area. For more information, see [Section 3.10, “Inventory Options,” on page 42](#).

To enable the *Inventory* button, you need to select credentials for the machines you’re adding. For more information, see [Section 3.8, “Inventory Credentials,” on page 40](#). The numbers that are displayed in the *Inventory* button indicate how many individual IP Addresses there are in the subnet you specified.

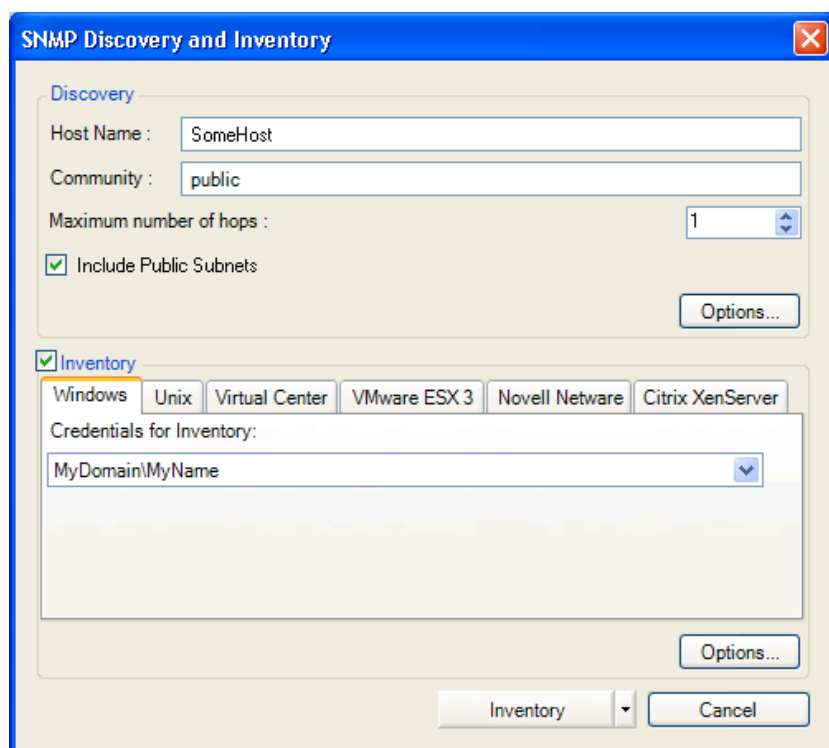
You can also start monitoring the machines when they are inventoried and/or add them to existing groups. For more information, see [Section 3.10, “Inventory Options,” on page 42](#).

3.7.7 SNMP

Used to discover machines by using Simple Network Management Protocol.

Type the *Host Name* and *Community*, then select the *Maximum number of hops*. Select whether to include *Public Subnets*.

Figure 3-6 *SNMP Discovery and Inventory Dialog Box*

The image shows a Windows-style dialog box titled "SNMP Discovery and Inventory". It has a blue title bar with a close button. The dialog is divided into two main sections. The top section, labeled "Discovery", contains fields for "Host Name" (with the value "SomeHost"), "Community" (with the value "public"), and "Maximum number of hops" (a spinner box set to "1"). There is a checked checkbox for "Include Public Subnets" and an "Options..." button. The bottom section, labeled "Inventory", has a checked checkbox and a tabbed interface with tabs for "Windows", "Unix", "Virtual Center", "VMware ESX 3", "Novell Netware", and "Citrix XenServer". The "Windows" tab is selected, showing a "Credentials for Inventory:" label and a dropdown menu with the value "MyDomain\MyName". There is another "Options..." button. At the bottom of the dialog are two buttons: "Inventory" and "Cancel".

NOTE: Clear the *Inventory* check box to run a discovery operation without inventorying server details.

Discovery options can be configured by clicking the *Options* button in the Discovery area. For more information, see [Section 3.10, “Inventory Options,” on page 42](#).

To enable the *Inventory* button, you need to select credentials for the machines you’re adding. For more information, see [Section 3.8, “Inventory Credentials,” on page 40](#).

You can also start monitoring the machines when they are inventoried and/or add them to existing groups. For more information, see [Section 3.10, “Inventory Options,” on page 42](#).

3.8 Inventory Credentials

To inventory discovered machines, select the *Inventory* check box in the Discovery and Inventory dialog boxes and then provide login credentials for Windows, UNIX, Virtual Center, VMware ESX, Novell NetWare, or Citrix XenServer machines, depending on the type of machines in use. For the Inventory From CSV File dialog box, the *Default Inventory Parameters* area is disabled if the CSV file contains credentials for its machines.

To provide inventory credentials, in the *Inventory* area, select credentials from the *Credentials for Inventory* drop-down list on the appropriate tab. For every machine in the list or range, PlateSpin Recon tries each type of credential provided until it finds one that matches. If the credentials are not valid, the machine is not inventoried. To learn more about adding credentials, refer to [Section 4.5, “Credentials Manager,” on page 69](#).

For inventorying UNIX or Citrix XenServer machines, select the *Use Sudo* check box to use superuser privileges for connections. Local policies can sometimes prevent you from establishing an SSH session as a `root` user. A user with sudo privileges can get around this restriction by establishing an SSH session as a non-`root` user and then using `sudo` to temporarily “switch” to `root` user privileges. To do this, create the credentials using that user’s username and password. PlateSpin Recon will use the `sudo` command to perform any operations that require `root` privileges.

For inventorying a Virtual Center, credentials must have a minimum read-only user status for the Virtual Center. Virtual Center automatically includes as users any user with Administrator privileges for the machine where Virtual Center is installed.

For inventorying Novell NetWare, NetWare Administrator credentials need to be specified. The credentials need permission to write to the Novell SYS: volume. As well, TCP port 524 needs to be open on the firewall. For more information see [Section 1.3, “Supported Platforms,” on page 13](#).

When running Inventory on NetWare machines, you also need to assign a set of Windows credentials which must be either:

- ♦ Part of the local Administrator group on the machine hosting the PlateSpin Recon server
- ♦ Have domain credentials associated with the following security policies:
 - ♦ Act as part of the operating system
 - ♦ Create a token object
 - ♦ Replace a process level token
 - ♦ Log on as a batch job
 - ♦ Log on as a service

3.9 Discovery Options

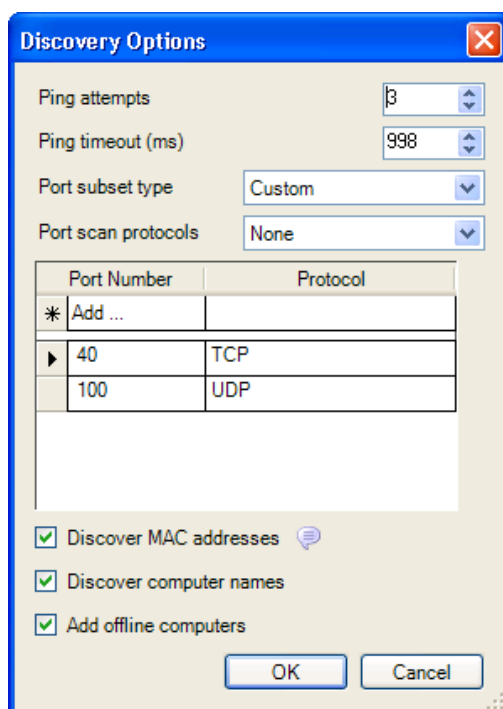
The following section discusses discovery options common to these dialog boxes:

- ♦ Domain Discovery and Inventory
- ♦ IP Range Discovery and Inventory
- ♦ Subnet Discovery and Inventory
- ♦ SNMP Discovery and Inventory

To configure discovery options:

- 1 Click the *Options* button in the *Discovery* area.
The Discovery Options dialog box is displayed.

Figure 3-7 *Discovery Options Dialog Box*



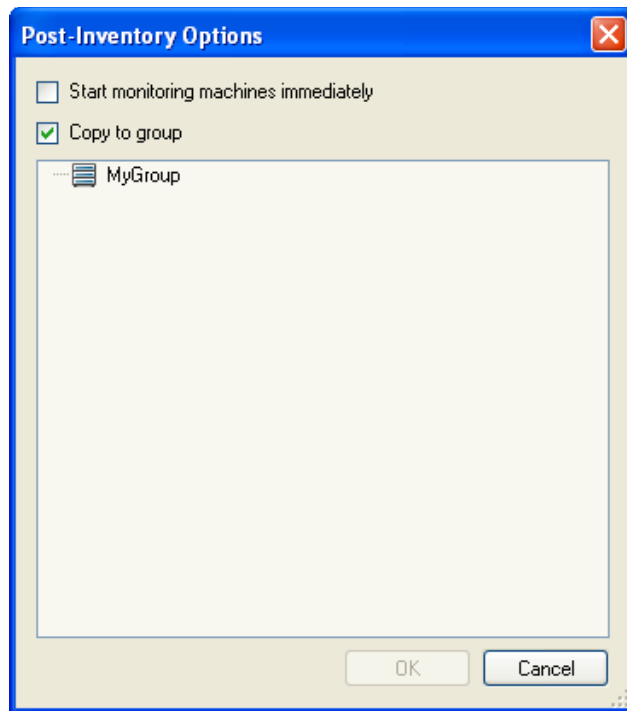
- 2 Configure options as desired.
- 3 If selecting *Custom* for the *Port subset type*, click *Add* to assign individual scan protocols to specific ports. Type the *Port Number* and select the *Protocol* from the drop-down cell. Hit *Enter* to apply your edits.
Most Common subset type refers to ports for common Windows network services. *All Standard* subset type refers to ports below 1025 used by standard protocols.
- 4 Select *Discover MAC addresses* to find MAC addresses of machines in the same ethernet segment as the PlateSpin Recon Server.
- 5 Select *Discover computer names* to include the computer names in the discovery poll.
- 6 Select *Add offline computers* to discover computers even if they are currently offline.
- 7 Click *OK*.

3.10 Inventory Options

These options apply if the *Inventory* check box is selected in *Discovery* and *Inventory* dialog boxes.

To configure inventory options:

- 1 Click the *Options* button in the *Inventory* area to configure post-inventory actions.



- 2 Select *Start monitoring machines immediately* to begin monitoring each machine discovered when the inventory is complete. Leave the check box cleared to start monitoring later. Select *Copy to group* and then select a target Group from those displayed in the list box to copy all the inventoried machines to that group.
- 3 Click *OK*.

3.11 Viewing Inventory Data and Monitoring Servers

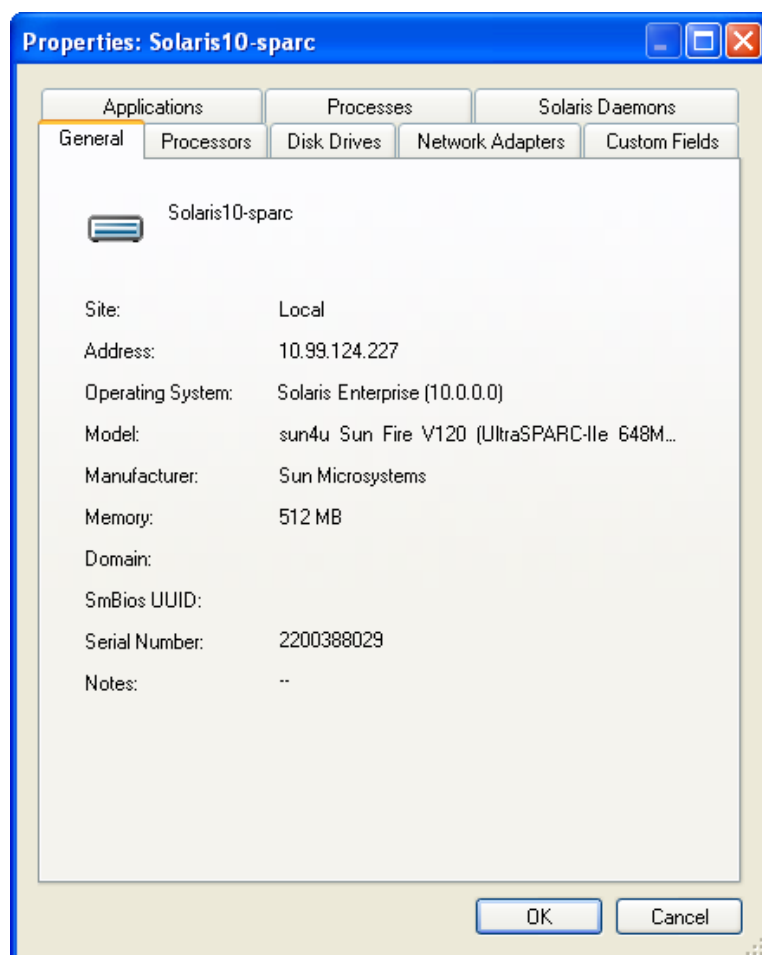
On completion of an inventory operation, you can view inventory data on servers that are listed in the Data Center Explorer tree. To view dynamic data, servers must be monitored for a period of time.

- ♦ [Section 3.11.1, “Viewing Properties,” on page 43](#)
- ♦ [Section 3.11.2, “Refreshing the Inventory,” on page 44](#)
- ♦ [Section 3.11.3, “Collecting Dynamic Data by Monitoring Servers,” on page 44](#)
- ♦ [Section 3.11.4, “Viewing Dynamic Data,” on page 45](#)
- ♦ [Section 3.11.5, “Viewing Logs,” on page 45](#)
- ♦ [Section 3.11.6, “Custom Fields,” on page 45](#)

3.11.1 Viewing Properties

Properties can assist with viewing static data for inventoried servers. To view inventory information, right-click a server and click *Properties*.

Figure 3-8 Sample View of the Properties Dialog Box



The tabs available in the Properties dialog box change based on the type of server being viewed.

3.11.2 Refreshing the Inventory

To refresh server inventory information in the Data Center Explorer, right-click a server or multiple servers and click *Refresh Inventory* on the menu. Inventory Refresh jobs are scheduled and run in the Job Explorer.

3.11.3 Collecting Dynamic Data by Monitoring Servers

- 1 Right-click the inventoried server and click *Start Monitoring*, or monitor multiple servers by selecting them or their parent group, filter, or site, then click *Start Monitoring*.

A job is created in the Job Explorer. For more details, refer to [Section 4.2, "Job Explorer,"](#) on [page 67](#).

- 2 After data is collected, click the server to view the corresponding chart in the Chart Viewer. With at least five minutes of data collected, the Chart Viewer can display workloads and utilization data.

To stop monitoring, right-click the monitored server and select *Stop Monitoring* . In the confirmation dialog box, click *Yes* to confirm or *No* to exit the dialog box and continue monitoring.

Although PlateSpin Recon might not be currently monitoring a server, it is still possible to view charts and reports showing the previously collected server data.

You can monitor multiple servers by selecting them individually or by using groups, filters, or sites.

3.11.4 Viewing Dynamic Data

You can view workloads and utilization data in a chart after collecting at least five minutes worth of data. Analysis Reports require a minimum of one hours worth of data. If there is insufficient data, the report is displayed without dynamic data. You can view collected data in a chart by selecting a monitored server.

See [Section 5.1, “Chart Viewer,” on page 75](#) and [Section 5.2, “Report Explorer,” on page 79](#) for more information.

3.11.5 Viewing Logs

Log files reflect the status of data collection for monitored servers. To inspect a log, right-click a server and click *Logs > View Logs*.

3.11.6 Custom Fields

Custom fields are available for entering server information that cannot be inventoried by PlateSpin Recon. These non-inventory fields can assist with identifying candidates for planning and consolidation. The information is stored in the visible table name.

To specify custom fields for multiple servers:

- 1 Select a server or multiple servers in the Data Center Explorer.
- 2 Right-click the selection and click *Properties* for a single server, or click *Set Custom fields* for multiple servers.

The Properties dialog box is displayed.

- 3 Click the *Custom Fields* tab.
- 4 Select a field to view its definition and parameters.

Figure 3-9 Custom Fields

Properties: DEV

General Processors Disk Drives Network Adapters Applications
Processes Microsoft Virtual Server Windows Services Custom Fields

Rack Units	2
Energy Consumption (W)	330
Heat Dissipation (BTU/hr)	1125
Normalized Speed	4.714753

Heat Dissipation (BTU/hr)
Amount of heat generated by a server. Heat Dissipation (BTU/hr) = Energy...

Reset Reference Values

Company	
Department	
Location City	
Location State	
Location Country	
Function	
Rack	
Billing Code	
Usage	
Lease Expires	
Warranty Expires	

Billing Code
Billing code of the server

OK Cancel

5 Type or select information in any of the fields, including those already filled with data.

6 Click *OK*.

If necessary, click *Reset Reference Values* to return all fields to their default values.

The database is updated for the selected servers. Report templates can filter for particular custom fields.

3.12 Working with Remote Data

Working with performance monitoring and resource utilization data external to the local instance of PlateSpin Recon allows you to take advantage of the PlateSpin Recon capacity planning, analysis, reporting and chargeback capabilities, regardless of the source of information. Data can be imported on a one-time or recurring basis.

3.13 Working with PlateSpin Recon Snapshots

PlateSpin Recon snapshots (*.prs files) contain data about machines, raters, report templates, generated reports, projects, scenarios, and server templates exported by other instances of PlateSpin Recon, or from the current instance as a backup.

You can also import snapshots exported from previous versions of PlateSpin Recon. In this case, a snapshot upgrade is performed before the data is brought into PlateSpin Recon. There are some things to watch for:

- ♦ Importing a large snapshot, for instance if the snapshot was compressed before export, requires two to five times the size of the old snapshot in disk space.
- ♦ Snapshots exported from PlateSpin Recon 3.0 and later are supported.
- ♦ Snapshots from PlateSpin Recon 3.1 are only supported if all hot fixes were applied to PlateSpin Recon before the snapshot was exported.

You can import a snapshot either into the default Local site, or into a new local site.

- ♦ [Section 3.13.1, “Exporting a PlateSpin Recon Snapshot,” on page 47](#)
- ♦ [Section 3.13.2, “Importing a PlateSpin Recon Snapshot,” on page 49](#)

3.13.1 Exporting a PlateSpin Recon Snapshot

- 1 Right-click in the Data Center Explorer and select *Sites > Export Snapshot*.

The Export Site Snapshot dialog box is displayed.

- 2 Select export objects to include in the snapshot by enabling the check box next to the objects.
When a report template is included in the snapshot, its generated reports can also be included. Similarly, if a project is included in the snapshot, its scenarios can be included.
- 3 Click *Next*.

If you selected to export *Machines* as part of the snapshot, the Step 1b dialog box is displayed.

Figure 3-10 Export Site Snapshot: Step 1b Dialog Box

Export Site Snapshot

Step 1b of 3: Specify the machines filter

Summary Data Granularity

☐ 5 Minute (Used in charts only)

☒ 1 Hour (Used in charts, reports, and projects)

Summary Data Range

☒ All Data

☐ Last Hour(s)

☐ Between and

☐ Monitor Logs

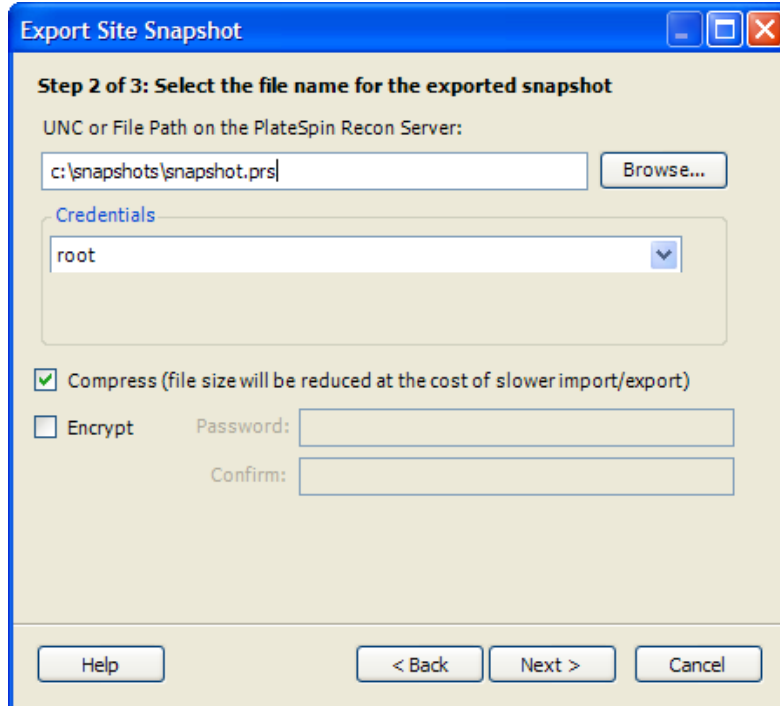
☐ Attached Credentials

Otherwise, proceed to [Step 2 on page 47](#).

- 4 Select summary options for the data to be exported, such as a data range time frame and whether to include Logs and/or Credentials.
- 5 Click *Next*.

The Step 2 dialog box is displayed.

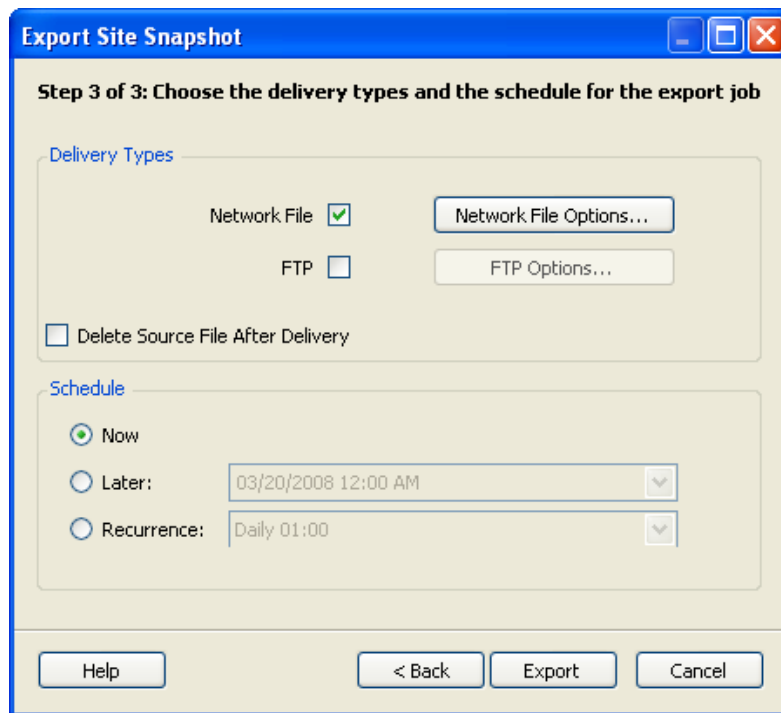
Figure 3-11 Export Site Snapshot: Step 2 Dialog Box



- 6 Enter the UNC (\\<server>\<path>\<filename>) or file path (C:\<folder>\<filename>) where you want to save the snapshot. You can also click *Browse* to select a location.
- 7 Select or create the appropriate credentials to access the location where you want to save the exported snapshot. You need write permission to this location.
- 8 Choose whether to *Compress* or *Encrypt* the file. For an encrypted snapshot file, specify and confirm a password for decryption.
- 9 Click *Next*.

The Step 3 dialog box is displayed.

Figure 3-12 *Export Site Snapshot: Step 3 Dialog Box*



- 10 If you want the snapshot delivered to a network and/or FTP location after it is exported, select a Delivery Type. To configure the delivery method options, click the option button next to the selected delivery method. If you are unsure of the delivery configurations, check with your System Administrator.

The file path must be accessible to the PlateSpin Recon Server.

You can specify whether to delete the source snapshot after it is delivered by enabling the *Delete Source File After Delivery* check box.

You can also specify whether you want the snapshot delivered now, later, or on a recurring schedule.

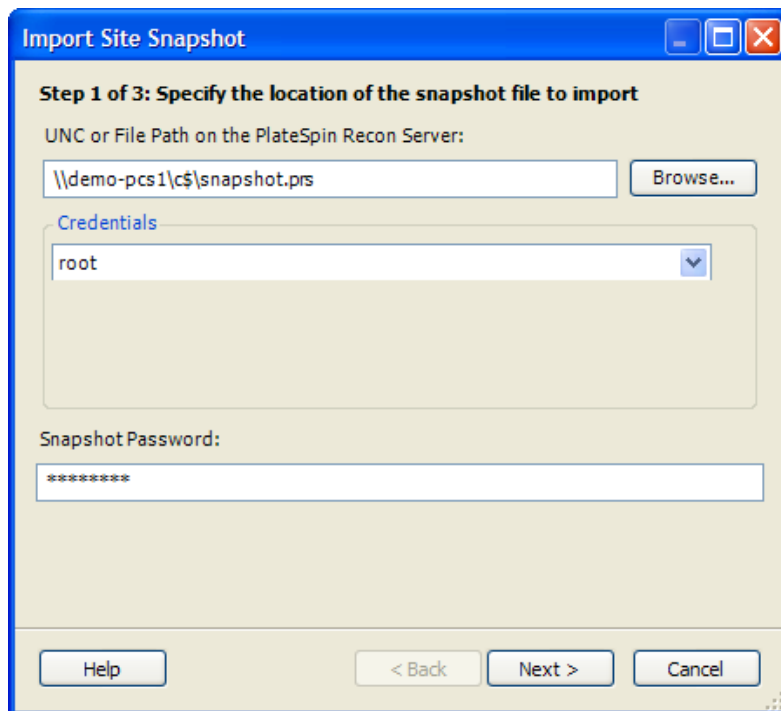
- 11 Click *Export*.

The data snapshot is ready for importing into another instance of PlateSpin Recon.

3.13.2 Importing a PlateSpin Recon Snapshot

- 1 Determine the type of local site (the default local site or a new local site) you want to import the required snapshot into. To import into a new local site, you need to create the site before importing the snapshot.
- 2 In the Data Center Explorer, right-click the target local site and select *Sites > Import Snapshot*.
The Import Site Snapshot dialog box is displayed.

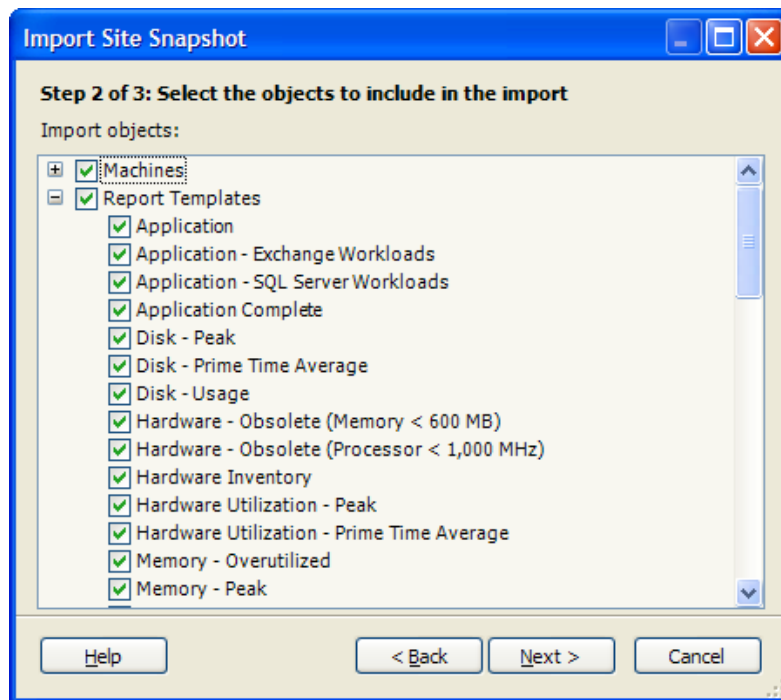
Figure 3-13 *Import Site Snapshot Dialog Box*

The image shows a Windows-style dialog box titled "Import Site Snapshot". It has a blue title bar with standard minimize, maximize, and close buttons. The main area is light beige. At the top, it says "Step 1 of 3: Specify the location of the snapshot file to import". Below this, it asks for the "UNC or File Path on the PlateSpin Recon Server:" and provides a text input field containing "\\demo-pcs1\c\$\snapshot.prs" and a "Browse..." button. Underneath is a section labeled "Credentials" with a dropdown menu showing "root". At the bottom of the main area, it asks for the "Snapshot Password:" with a text input field containing seven asterisks. The bottom of the dialog box contains four buttons: "Help", "< Back", "Next >", and "Cancel".

- 3** Type the *UNC or Local File Path* of where the snapshot you want to import is located, or click *Browse* to find it.
- 4** Select or add the *Credentials* needed to access the location.
- 5** If the file was encrypted on export, type the *Snapshot Password* to decrypt it.
- 6** Click *Next*.

The Step 2 dialog box is displayed..

Figure 3-14 Import Site Snapshot: Step 2 Dialog Box

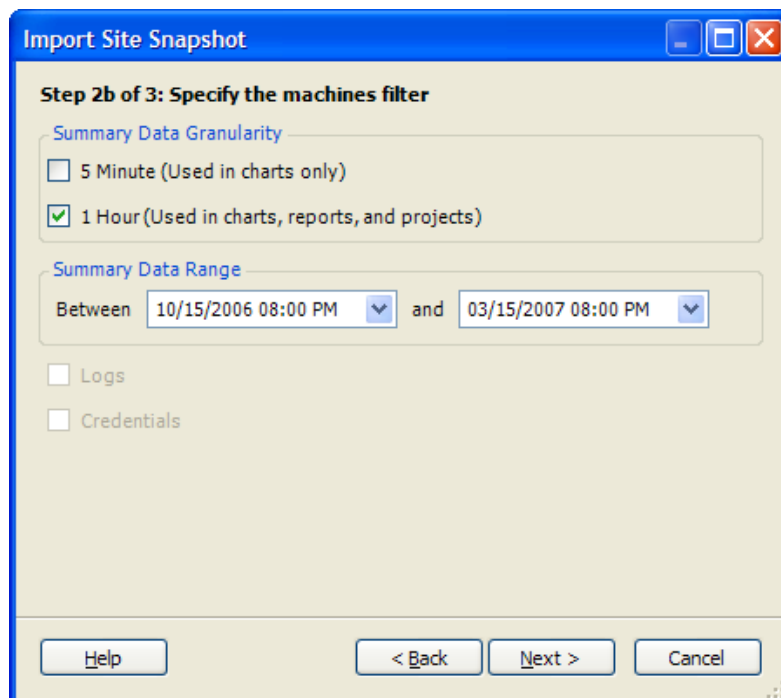


7 Enable the check boxes next to the *Import objects* you want to import.

8 Click *Next*.

If the imported objects include *Machines* with monitoring data, the Step 2b dialog box is displayed. Otherwise, proceed to [Step 2 on page 49](#).

Figure 3-15 Import Site Snapshot: Step 2b Dialog Box



- 9 Specify the data options for the machines filter.
- 10 Click *Next*.
The Step 3 dialog box is displayed.
- 11 Use the drop-down list to select the target site for the imported data.
- 12 Click *Import*.

After the data is imported, work with it using the Data Center Explorer under the *Sites* node.

3.14 Working with Data Collected by Other PlateSpin Recon Instances

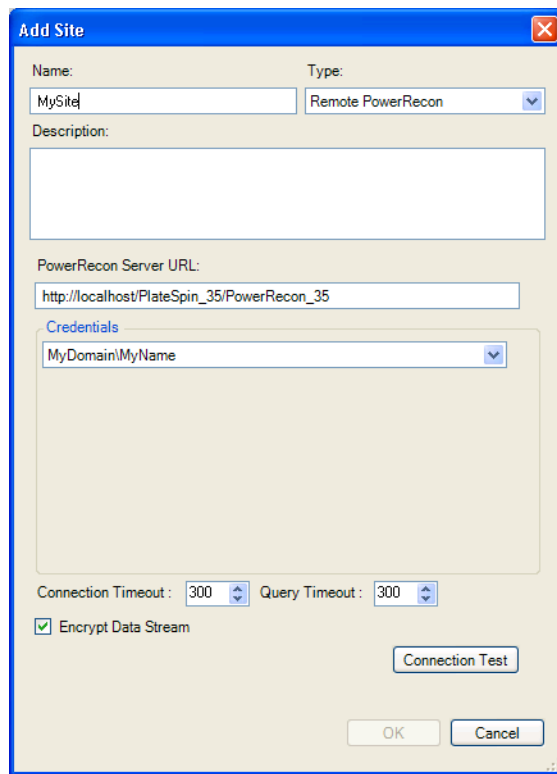
You can use performance and utilization data collected by other instances of PlateSpin Recon, such as in a setup where multiple PlateSpin Recon instances handle monitoring and data collection, with one acting as a central PlateSpin Recon instance for generating reports and developing consolidation plans.

NOTE: Machines imported from a remote PlateSpin Recon site can be viewed in reports and charts once the data has been synchronized, as described below. Further inventorying or monitoring cannot take place, though. To do this, perform the actions at the original remote PlateSpin Recon site and then import the machines again. You will also need to synchronize the data again.

To work with data collected by another instance of PlateSpin Recon:

- 1 Create a new Remote PlateSpin Recon site to connect to the remote PlateSpin Recon instance:
 - 1a Right-click in the Data Center Explorer and select *Sites > Add Site*.
The Add Site dialog box is displayed.
 - 1b Select *Remote PlateSpin Recon* as the new site *Type* and specify the required basic parameters for the new site, including a *Name* and an optional *Description*:

Figure 3-16 Add Site: Remote PlateSpin Recon Dialog Box

The image shows a Windows-style dialog box titled "Add Site". It has a blue title bar with a close button (X) in the top right corner. The dialog is divided into several sections. The top section has two labels: "Name:" and "Type:". The "Name:" label is followed by a text box containing "MySite". The "Type:" label is followed by a dropdown menu showing "Remote PowerRecon". Below this is a "Description:" label followed by a large empty text area. The next section is labeled "PowerRecon Server URL:" and contains a text box with the URL "http://localhost/PlateSpin_35/PowerRecon_35". Below that is a "Credentials" section with a dropdown menu showing "MyDomain\MyName". At the bottom, there are two spin boxes for "Connection Timeout:" and "Query Timeout:", both set to "300". Below these is a checked checkbox labeled "Encrypt Data Stream". At the very bottom, there are three buttons: "Connection Test", "OK", and "Cancel".

1c Type the *PlateSpin Recon Server URL* in the format:

`http://<host>[:port]/PlateSpin_35/PlateSpin Recon_35`

1d From the *Credentials* drop-down list, select the appropriate credentials or select *<New Credential>* and add new credentials.

1e If you want the data stream encrypted, ensure that the *Encrypt Data Stream* check box is selected.

1f (Optional) Click *Connection Test* to verify connectivity with the remote server.

1g Click *OK*.

Note the new site under *Sites* in the Data Center Explorer.

2 Right-click the Remote PlateSpin Recon site you just created, then select *Sites > Import Machines* from the menu.

3 Right-click the Remote PlateSpin Recon site again, then select *Sites > Synchronize Data* from the menu.

3.15 Working with Data Collected by Other Performance Monitoring Tools

You can import and work with performance and utilization data collected by other performance monitoring software, such as Microsoft Operations Manager (MOM) 2005, Microsoft Windows Performance Monitor (`perfmon.exe`), and HP Reporter. You can also use the PlateSpin Recon Performance Data Import feature.

NOTE: Machines imported can be viewed in reports and charts once the data has been synchronized. Further inventorying or monitoring cannot take place, though. To do this, perform the actions at the original remote site and then import the machines again. You will also need to synchronize the data again.

PlateSpin Recon supports summarization of collected data within 5-minute and 60-minute intervals, regardless of how the selected connector summarizes and collects data. PlateSpin Recon requires that data collected by connectors be summarized over intervals of 60 minutes or less. We recommend importing connector data summarized over 60-minute intervals, especially for systems whose utilization data covers a period of longer than one month.

For PlateSpin Recon to generate reliable consolidation plans, we recommend that you collect utilization data for a period of one to three months before importing it into PlateSpin Recon.

For the capability to work with data collected by third-party software, you must:

- ♦ Create a new Remote site of the appropriate type and select an appropriate connector.
- ♦ Inventory the remote machines.
- ♦ Synchronize data in the new site with that collected by the remote monitoring tool.
- ♦ [Section 3.15.1, “Required Performance Counters,” on page 54](#)

3.15.1 Required Performance Counters

You must ensure that the required performance and utilization data is actually gathered for PlateSpin Recon counters in the required categories.

Table 3-8 *Required PlateSpin Recon Counters*

PlateSpin Recon Counter	Expressed in
Processor utilization	Percentage of total processor time
Memory utilization	Total number of free bytes
Disk IO	Bytes per second
Network IF I/O	Bytes per second

NOTE: NetWare machines only require two counters (CPU and disk) to be gathered for successful consolidation. For more information see [Section 1.3, “Supported Platforms,” on page 13](#).

Where required, PlateSpin Recon provides mechanisms for converting third-party performance and utilization data to the required PlateSpin Recon format.

- ♦ [“Working with Microsoft Operations Manager \(MOM\) 2005 Data” on page 55](#)
- ♦ [“Working with Windows PerfMon Counter Logs” on page 56](#)

Working with Microsoft Operations Manager (MOM) 2005 Data

- 1 Download and install the PlateSpin Management Pack for MOM 2005 (see Knowledge Base Article Q20903 (<http://support.platespin.com/kb2/article.aspx?id=20903>)) and collect data by using MOM.
- 2 Create a new Remote site for connecting to MOM:
 - 2a Right-click in the Data Center Explorer and select *Sites > Add Site*.
The Add Site dialog box is displayed.
 - 2b Select *Remote* as the *Type* of the new site and specify the required basic parameters for the new site, including a *Name* and an optional *Description*.

Figure 3-17 Add Remote Site

The screenshot shows the 'Add Site' dialog box. The 'Name' field is 'MySite' and the 'Type' dropdown is set to 'Remote'. The 'Description' field is empty. The 'Connector Type' dropdown is set to 'Microsoft Operations Manager 2005'. Under the 'Database Connection' section, 'Database Hostname' is 'MyDB', 'Database Instance' is 'OnePoint', 'Connection Timeout' is 300, and 'Query Timeout' is 300. Under the 'Database Credentials' section, the dropdown shows 'MyDomain\MyName'. At the bottom right, there is a 'Test Settings' button, and at the bottom center, there are 'OK' and 'Cancel' buttons.

- 2c From the *Connector Type* drop-down list, select *Microsoft Operations Manager 2005*.
- 2d Type the *Database Hostname* for the remote server. *Database Instance* defaults to *OnePoint* and must be changed to reflect the Microsoft Operation Manager 2005 database instance in use.

- 2e** Select the appropriate credentials or select *<New Credential>* to add new credentials. To verify connectivity with the remote database, click *Test Settings*.
- 2f** Click *OK*.

The new site is added under *Sites* in the Data Center Explorer.
- 3** Inventory the machines:
 - 3a** In the Data Center Explorer, right-click the remote site and select *Sites > Inventory Machines*.

The Inventory Machines dialog box is displayed.
 - 3b** Select the check boxes next to the machines you want to inventory, specify the required inventory parameters, then click *OK*.
- 4** Synchronize site data:
 - 4a** In the Data Center Explorer, right-click the required site and select *Sites > Synchronize Data*.
 - 4b** Select the check boxes next to the required machines. To select all machines in a site, select the check box next to the site.
 - 4c** Select the required data range, scheduling, and recurrence options.
 - 4d** Click *OK*.

Working with Windows PerfMon Counter Logs

- 1** Set up your Windows Performance Monitor (`perfmon.exe`) to gather counter data required by PlateSpin Recon. See [Section 3.15.1, “Required Performance Counters,” on page 54](#).
- 2** Export the data to a CSV file. For details, see your Windows documentation.
- 3** Create a new Remote site for working with Windows Performance Monitor logs:
 - 3a** Right-click in the Data Center Explorer and select *Sites > Add Site*.

The Add Site dialog box is displayed.
 - 3b** Select *Remote* as the *Type* of the new site and specify the required basic parameters for the new site, including a *Name* and an optional *Description*.

Figure 3-18 *Add Site: Windows Performance Monitor*

The screenshot shows the 'Add Site' dialog box. The 'Name' field contains 'MySite'. The 'Type' dropdown menu is set to 'Remote'. The 'Description' field is empty. The 'Connector Type' dropdown menu is set to 'Performance Monitor'. The 'UNC or File Path on the PowerRecon Server' field contains 'C:\PerfMonData\Import\PerfMonSite\11-2008.csv', and there is a 'Browse...' button next to it. The 'Credentials' section has a dropdown menu set to 'MyDomain\MyName'. At the bottom, there is a 'Test Settings' button, and 'OK' and 'Cancel' buttons.

3c From the *Connector Type* drop-down list, select *Performance Monitor*.

3d Type the path to (or browse to and select) the CSV file that contains data exported from Performance Monitor.

3e Select the required credentials or select *<New Credential>* to add new credentials.

3f Click *Test Settings* to verify UNC or local file path access.

3g Click *OK*.

The new site is added under *Sites* in the Data Center Explorer.

4 Inventory the machines:

4a In the Data Center Explorer, right-click the remote site and select *Sites > Inventory Machines*.

The Inventory Machines dialog box is displayed.

- 4b** Select the check boxes next to the machines you want to inventory, specify the required inventory parameters, then click *OK*.
- 5** Synchronize site data:
 - 5a** In the Data Center Explorer, right-click the required site and select *Sites > Synchronize Data*.
 - 5b** Select the check boxes next to the required machines. To select all machines in a site, select the check box next to the site.
 - 5c** Select the required data range, scheduling, and recurrence options.
 - 5d** Click *OK*.

3.16 Using HP Reporter Data with PlateSpin Recon

PlateSpin Recon features a specialized HP Reporter connector. It enables PlateSpin Recon to connect to an HP Reporter 3.8 database to import performance and resource utilization data for charts, reports and consolidation planning.

You can use any HP Reporter-collected data to generate PlateSpin Recon charts and reports. However, for consolidation planning you must configure Reporter to collect specific utilization data that correlates to data collected by specific PlateSpin Recon counters. For details, see [Section 3.15.1, “Required Performance Counters,” on page 54](#).

Because HP Reporter doesn’t collect data in exactly the same format as PlateSpin Recon, you need to specify appropriate formulas to transform the data from Reporter-collected metrics to PlateSpin Recon counters.

- ♦ [Section 3.16.1, “Using Formulas to Convert Metrics to Counters,” on page 58](#)
- ♦ [Section 3.16.2, “Example Formulas,” on page 59](#)

The HP Reporter data summarization interval is specified during the creation of the HP Reporter Metric List.

3.16.1 Using Formulas to Convert Metrics to Counters

You can use a formula to provide the following information to the PlateSpin Recon HP Reporter Connector:

- ♦ Which Reporter metric list contains metrics required by PlateSpin Recon
- ♦ How to convert data from measurement units used by Reporter units to those required by PlateSpin Recon

A valid HP Reporter Connector formula consists of the name of the required HP Reporter metric list, followed by a dot character, followed by an arithmetic expression that converts the metric into a PlateSpin Recon counter. Arithmetic expressions can contain any of the collected Reporter metrics, numeric constants, and arithmetic operations.

NOTE: Only metrics collected by the same metric list can be used in one formula, although different formulas can use different metric lists.

For example, PlateSpin Recon requires memory utilization data expressed in the number of free bytes available to the operating system, while Reporter collects the same data by using two separate metrics:

- ♦ Percentage of memory utilized (GBL_MEM_UTIL)
- ♦ Total size of the memory in gigabytes (GBL_MEM_PHYS)

In this case, the appropriate formula for converting collected metric data into the PlateSpin Recon counter format would be:

$$\text{GBL_MEM_PHYS} * (100 - \text{GBL_MEM_UTIL}) * 1024 * 1024 * 1024 / 100$$

3.16.2 Example Formulas

The following are examples of HP Reporter metrics that you can collect using CODA agents, along with corresponding conversion formulas. Note that To do consolidation planning, PlateSpin Recon requires data for all four counters. Because HP Reporter can collect metric data with several different agents, and the same PlateSpin Recon counters can be calculated using different HP Reporter metrics, this list is not exhaustive and serves only as an example.

Table 3-9 Example HP Reporter Metrics

PlateSpin Recon Counter	HP Reporter Metrics	Conversion Formula
Processor Utilization	GBL_CPU_TOTAL_UTIL	<MetricList>.GBL_CPU_TOTAL_UTIL
	BYCPU_CPU_TOTAL_UTIL	<MetricList>.BYCPU_CPU_TOTAL_UTIL
Memory Utilization	GBL_MEM_PHYS and GBL_MEM_UTIL	<MetricList>.(GBL_MEM_PHYS*(100-GBL_MEM_UTIL))*1024*1024*1024/100
	GBL_MEM_FREE	<MetricList>.(GBL_MEM_FREE)*1024*1024
Disk Utilization	GBL_DISK_PHYS_BYTE_RATE	<MetricList>.GBL_DISK_PHYS_BYTE_RATE
	BYDSK_PHYS_READ_BYTE_RATE	<MetricList>.BYDSK_PHYS_READ_BYTE_RATE
Network Utilization	BYNETIF_OUT_BYTE_RATE and BYNETIF_IN_BYTE_RATE	<MetricList>.BYNETIF_OUT_BYTE_RATE + BYNETIF_IN_BYTE_RATE

Data for certain HP Reporter metrics and PlateSpin Recon counters is collected separately for each device instance. For example, the BYCPU_CPU_TOTAL_UTIL metric gives a separate utilization figure for each CPU in the system. To distinguish among different device instances, include a metric with a device identification string in the list when you set up the HP Reporter metric lists.

The PlateSpin Recon HP Reporter connector identifies the following metrics as containing device instance names:

Table 3-10 *PlateSpin Recon HP Reporter Metrics*

Functional Area	Metric
Processors	BYCPU_ID
Disks	BYDSK_DEVNAME or BYDSK_ID
Network Interface Cards	BYNETIF_NAME or BYNETIF_ID
File Systems	FS_DIRNAME, FS_DEVNAME or FS_DEVNO

To work with HP Reporter data:

- 1** Configure your HP Reporter to collect data in the required performance and utilization measurement categories. See [Section 3.15.1, “Required Performance Counters,” on page 54](#).
- 2** (Conditional) If you are using Oracle with HP Reporter, install Oracle client software on the PlateSpin Recon Server host. For details, see your Oracle documentation.
- 3** Create a new Remote site for connecting to the HP Reporter database:
 - 3a** Right-click in the Data Center Explorer and select *Sites > Add Site*.
The Add Site dialog box is displayed.
 - 3b** Specify a *Name* and, optionally, a *Description*.
 - 3c** Click the *Type* drop-down and select *Remote* from the list.

Figure 3-19 Add Remote Site Dialog Box

The image shows a Windows-style dialog box titled "Add Site". It has a blue title bar with a close button (X) in the top right corner. The dialog is divided into several sections. At the top, there are two labels: "Name:" and "Type:". The "Name:" label is followed by a text input field containing "MySite". The "Type:" label is followed by a dropdown menu showing "Remote". Below these is a "Description:" label followed by a large, empty text area. Underneath the description area is a "Connector Type:" label followed by a dropdown menu showing "HP Reporter". Below the connector type is a tabbed interface with two tabs: "Settings" (which is selected and highlighted with an orange border) and "Counters". The "Settings" tab contains a section titled "Database Connection" with several fields: "Server Name:" with input "MyServer", "Database Name:" with input "reporter", "Connection Timeout:" with a spinner set to "300", "Query Timeout:" with a spinner set to "300", "Database Type:" with a dropdown showing "SQL Server 2005", and "Database Port:" with input "1433". Below the "Database Connection" section is another section titled "Database Credentials" with a dropdown menu showing "root". At the bottom right of the "Settings" tab is a button labeled "Test Settings" with a yellow warning icon. At the very bottom of the dialog are two buttons: "OK" and "Cancel".

- 3d** From the *Connector Type* drop-down list, select *HP Reporter*.
HP Reporter-related *Settings* and *Counters* fields are displayed.
- 3e** On the *Settings* tab, provide the required information for connecting to the HP Reporter database: database hostname, database or database instance name, database type, port number (if a non-default port number is used), and valid credentials for accessing the database.
- When you synchronize a large dataset from HP Reporter, you might need to set the *Connection timeout* and/or the *Query timeout* values to 0, which means infinite/unlimited timeout.
- 3f** On the *Counters* tab, for the four required performance and utilization data categories, provide counter formulas for the system to use when converting HP Reporter metrics to PlateSpin Recon counters. See [Section 3.16.1, “Using Formulas to Convert Metrics to Counters,”](#) on page 58.
- 3g** Click *OK*.
The new site is added under *Sites* in the Data Center Explorer.

- 4 Inventory the machines:
 - 4a In the Data Center Explorer, right-click the remote site you just created and select *Sites > Inventory Machines*.

The Inventory Machines in Site <sitename> dialog box is displayed.
 - 4b Select the check boxes next to the machines you want to inventory.
 - 4c Select the required credentials or select <New Credential> to add new credentials.
 - 4d Click *OK*.
- 5 Synchronize site data:
 - 5a In the Data Center Explorer, right-click the required site and select *Sites > Synchronize Data*.
 - 5b Select the check boxes next to the required machines. To select all machines in a site, select the check box next to the site.
 - 5c Select required data range, scheduling, and recurrence options.
 - 5d Click *OK*.

For in-depth information about using PlateSpin Recon in conjunction with HP Reporter, see Knowledge Base article [Q21088 \(http://support.platespin.com/kb2/article.aspx?id=21088\)](http://support.platespin.com/kb2/article.aspx?id=21088).

3.17 Using Performance Data Import to Import Generic Data into PlateSpin Recon

PlateSpin has developed a robust Performance Data Import connector for PlateSpin Recon that allows users to import collected performance data from many third-party tools on the market into PlateSpin Recon to take advantage of its charting, reporting, and consolidation planning features.

When collected data has been moved from its proprietary location and format into a set of tables in the PlateSpin Recon monitoring database, PlateSpin Recon can import the data and make use of it like any other performance data. You can also set up a scheduled or recurring import, so the PlateSpin Recon analysis doesn't have to be a one-time thing.

Almost all performance data gathering tools use a different method and format for storing the gathered data. To enable PlateSpin Recon to successfully import any data, regardless of its original format, you need to use an Extract, Transform, and Load (ETL) tool. This type of tool takes the data, reshapes it, and then loads it into the tables where PlateSpin Recon expects to find it. See [Section 3.17.2, "Preparing Data for Import into PlateSpin Recon," on page 64](#).

The workloads in the data being imported need to be of the type that PlateSpin Recon can inventory under normal circumstances. All necessary ports between PlateSpin Recon and the workload being imported must be open.

As with other remote data connectors, data to be imported should contain one or more data points for every hour. Also, to use Consolidation Planning, specific counters need to exist in the data. See [Section 3.15.1, "Required Performance Counters," on page 54](#).

- ♦ [Section 3.17.1, "Performance Data Import Tables Data Structure," on page 63](#)
- ♦ [Section 3.17.2, "Preparing Data for Import into PlateSpin Recon," on page 64](#)

3.17.1 Performance Data Import Tables Data Structure

There are four tables used for importing of generic performance data:

- ♦ **ImportComputer table:** Contains information about computers that are to be imported into PlateSpin Recon.
- ♦ **ImportCounter table:** Contains information about all counters that were collected for any of the computers in the *ImportComputer* table.
- ♦ **ImportCounterInstance table:** Contains information about named instances of each counter for each computer. The combination of *Computer*, *Counter* and *Instance Name* must be unique.
- ♦ **ImportData table:** Contains monitoring data to be imported.

Each table contains one ID field, which acts as a primary key and is automatically generated by the database server when a new row is inserted into the table.

Table 3-11 *ImportComputer Table*

Column Name	Type	Description
ID	Integer	Generated automatically.
Name	varchar(255)	Network name of computer. This is the address or name of the computer. It allows PlateSpin Recon to address the computer over the network.

Table 3-12 *ImportCounter Table*

Column Name	Type	Description
ID	Integer	Generated automatically.
Name	varchar(255)	Unique name of the counter. This is the name of the counter as it will be shown in PlateSpin Recon charts. Some predefined counter names exist. See Section 3.15.1, "Required Performance Counters," on page 54.

Table 3-13 *ImportCounterInstance Table*

Column Name	Type	Description
ID	Integer	Generated automatically.
Computer	Integer	Reference to the ImportComputer table.
Counter	Integer	Reference to the ImportCounter table.
InstanceName	Integer	Name of the given counter instance on the specific computer.

Table 3-14 *ImportData Table*

Column Name	Type	Description
ID	Integer	Generated automatically.
Instance	Integer	Tells PlateSpin Recon for which computer, counter and counter instance this data value was collected.
ObservationDate	DateTime	Time the data was collected.
ObservationValue	float	Value of the counter.

3.17.2 Preparing Data for Import into PlateSpin Recon

To migrate data from a third-party location and format requires an Extract, Transform and Load (ETL) tool. There are many such tools on the market:

Open Source: CloverETL, Enhydra* Octopus, Pentaho* Data Integration, and Talend Open Studio.

You can also write proprietary data manipulation scripts, such as SQL queries, although the specifics of these tasks are beyond the scope of this document. For third-party tools, see the documentation that accompanies those tools. For scripts, refer to your in-house reference material on the languages and requirements.

Commercial: Informatica* PowerCenter*, Inaplex* Inaport*, and Altova* MapForce*.

Regardless of the tool you choose, the data transformation workflows are the same:

- ♦ Retrieve hostnames of the workloads you want to import and write them to the ImportComputer table
- ♦ Write the names of the required counters to the ImportCounter table
- ♦ Retrieve and write all performance data related to the above counters

To work with generic imported data:

- 1** Collect data in the required performance and utilization measurement categories. See [Section 3.15.1, “Required Performance Counters,” on page 54.](#)

- 2** Extract, transform, and load the collected performance data into the PlateSpin Recon Performance Data Import tables.

See [Section 3.17, “Using Performance Data Import to Import Generic Data into PlateSpin Recon,” on page 62](#), [Section 3.17.1, “Performance Data Import Tables Data Structure,” on page 63](#) and [Section 3.17.2, “Preparing Data for Import into PlateSpin Recon,” on page 64.](#)

- 3** Create a new Remote site for connecting to the Performance Data Import tables:

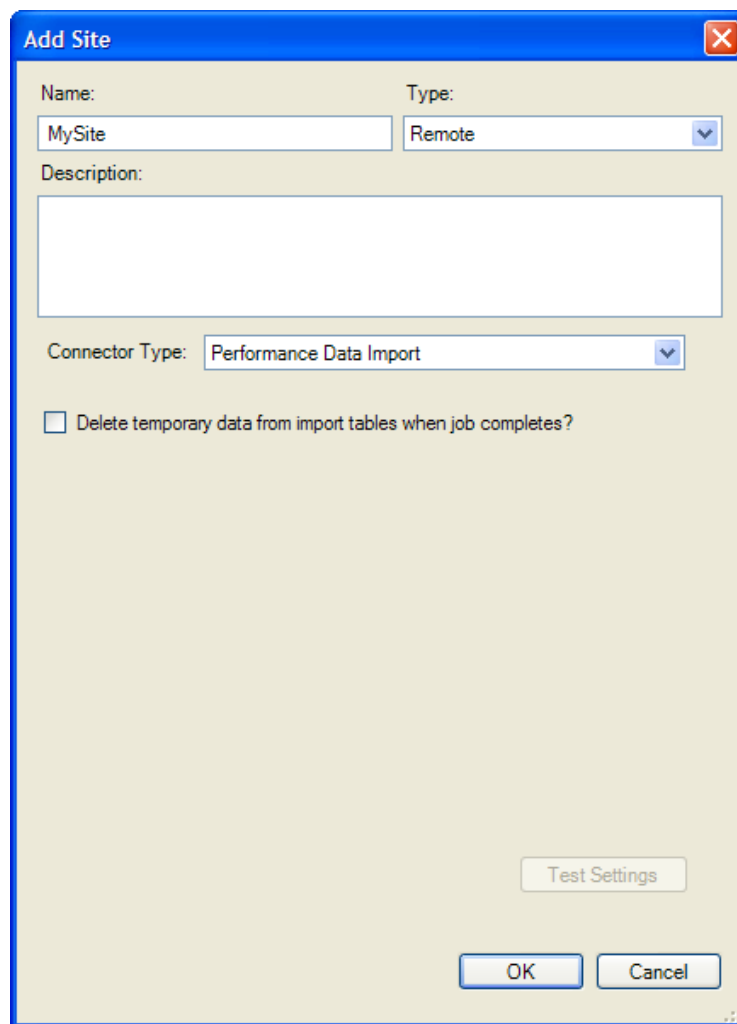
- 3a** Right-click in the Data Center Explorer and select *Sites > Add Site*.

The Add Site dialog box is displayed.

- 3b** Specify a *Name* and, optionally, a *Description*.

- 3c** Click the *Type* drop-down and select *Remote* from the list:

Figure 3-20 Sample Performance Data Import Tables Remote Site Dialog Box



The dialog box is titled "Add Site" and has a standard Windows-style title bar with a close button. It contains the following fields and controls:

- Name:** A text input field containing "MySite".
- Type:** A dropdown menu currently set to "Remote".
- Description:** A large, empty text area.
- Connector Type:** A dropdown menu currently set to "Performance Data Import".
- Delete temporary data from import tables when job completes?:** An unchecked checkbox.
- Test Settings:** A button located below the checkbox.
- OK** and **Cancel** buttons at the bottom right.

- 3d** From the *Connector Type* drop-down list, select *Performance Data Import*
The *Delete data from tables when done?* check box is displayed.
- 3e** To empty the Performance Data Import tables of the data loaded in **Step 2**, select the *Delete data from tables when done?* check box.
If you choose not to allow PlateSpin Recon to empty these tables, you must do this manually before performing another import.
- 3f** Click *OK*
The new site is added under *Sites* in the Data Center Explorer.
- 4** Inventory the machines:
- 4a** In the Data Center Explorer, right-click the remote site you just created and select *Sites > Inventory Machines*
The Inventory Machines in Site <sitename> dialog box is displayed.
- 4b** Select the check boxes next to the machines you want to inventory.

- 4c** Select the required credentials or select *<New Credential>* to add new credentials.
- 4d** Click *OK*.
- 5** Synchronize site data:
 - 5a** In the Data Center Explorer, right-click the required site and select *Sites > Synchronize Data*.
 - 5b** Select the check boxes next to the required machines. To select all machines in a site, select the check box next to the site.
 - 5c** Select required data range, scheduling, and recurrence options.

To perform a recurring import, you need to also set up a matching recurring Extract, Transform, and Load process to refresh the data in the PlateSpin Recon Performance Data Import tables. See your software vendor documentation for more details.
 - 5d** Click *OK*.

PlateSpin Recon Server and Client Tools

This section provides information about PlateSpin Recon utilities which enable you to better manage login credentials, analyze data center assets, and monitor operations.

- ♦ Section 4.1, “Server and Client Options,” on page 67
- ♦ Section 4.2, “Job Explorer,” on page 67
- ♦ Section 4.3, “Scheduled Job Explorer,” on page 68
- ♦ Section 4.4, “Log Explorer,” on page 69
- ♦ Section 4.5, “Credentials Manager,” on page 69
- ♦ Section 4.6, “PlateSpin Analyzer,” on page 71

4.1 Server and Client Options

Server and Client settings are available in the Options dialog box (*Tools > Options*). It is highly recommended that you leave the Server settings as is and *ONLY* change their values when directed to do so by PlateSpin Technical Support.

Some Client settings can be configured. For instance, it might be necessary to specify a path for various external applications or to change the option to check for a software update.

To return Server and Client options to their default values, click *Reset All*.


4.2 Job Explorer

Whenever a task is performed in PlateSpin Recon, a job is created for the task. Information on running, paused or completed jobs is shown in the Job Explorer, located at the bottom of the PlateSpin Recon window, by default.

If a job consists of several tasks, you can expand the job to see all of its tasks and the status of each task by clicking the plus sign next to the job description.

You can also see the steps that make up each task by selecting the task. The steps are displayed in the panel on the right of the Job Explorer.






Figure 4-1 Task steps displayed in the Job Explorer



Messages about the progress of a job are captured in a log file. Right-click any running job entry to display a pop-up menu, or use the toolbar buttons, and you can *Suspend*, *Resume*, or *Abort* any running job. The *Archive* command moves jobs to the Archived Jobs list. These jobs can be subsequently viewed by selecting *Tools > Archived Jobs*. You can view the step logs by right-clicking on a step and selecting *View Logs*. The Log Viewer dialog is displayed.

NOTE: The default setting is to purge logs and jobs every 14 days. To change this setting, select *Tools > Options*. The Options dialog box is displayed. Double-click *Advanced* under the *Maintenance* category on the Server page.

Table 4-1 Job Explorer commands

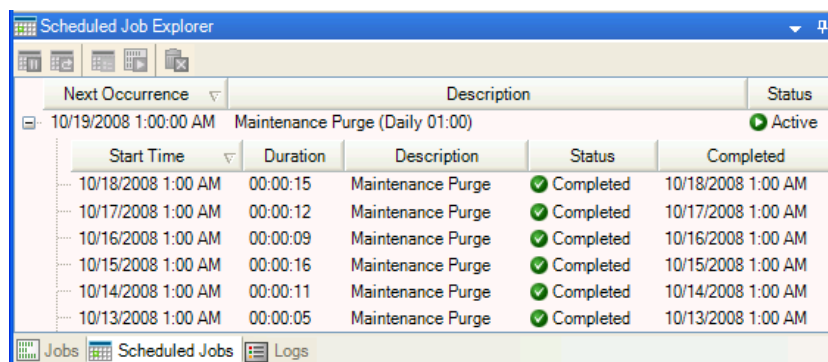
Commands	Tooltip/Description
	Suspend: Suspends a running job.
	Resume: Resumes a paused job.
	Abort: Abort the currently running job.
	Archive: Archives the selected job.
	View Logs: Opens a log viewer for the selected step.

4.3 Scheduled Job Explorer

A scheduled job is an activity set to occur at a given time and frequency. For example, it is possible to generate a Network Usage Report every morning at 9:00 AM for the next five days.

A scheduled job is comprised of two components: the action to be performed (the job), and the schedule for performing the activity. The schedule is a reusable component that can be shared amongst multiple jobs, such as creating three reports at 9:00 AM. The Scheduled Job Explorer can be used to view the status of the completed jobs.

Figure 4-2 Job status displayed in the Scheduled Job Explorer








Next Occurrence	Description	Status
10/19/2008 1:00:00 AM	Maintenance Purge (Daily 01:00)	Active

Start Time	Duration	Description	Status	Completed
10/18/2008 1:00 AM	00:00:15	Maintenance Purge	Completed	10/18/2008 1:00 AM
10/17/2008 1:00 AM	00:00:12	Maintenance Purge	Completed	10/17/2008 1:00 AM
10/16/2008 1:00 AM	00:00:09	Maintenance Purge	Completed	10/16/2008 1:00 AM
10/15/2008 1:00 AM	00:00:16	Maintenance Purge	Completed	10/15/2008 1:00 AM
10/14/2008 1:00 AM	00:00:11	Maintenance Purge	Completed	10/14/2008 1:00 AM
10/13/2008 1:00 AM	00:00:05	Maintenance Purge	Completed	10/13/2008 1:00 AM

Errors or warnings, if any, are also displayed. In addition, scheduled jobs can be run manually, or can have their recurrence changed.

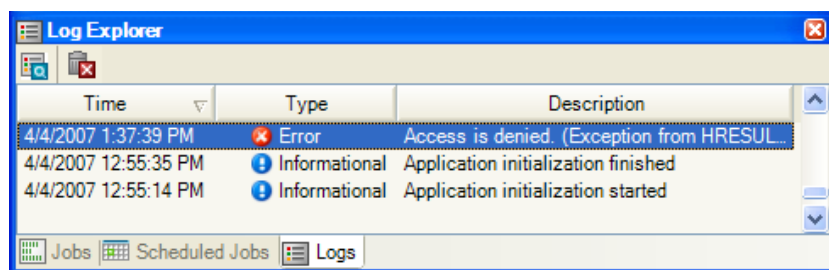
Table 4-2 *Scheduled Job Explorer commands*

Commands	Tooltip/Description
	Pause: Pauses an active or scheduled job.
	Resume: Resumes a paused job.
	Edit Recurrence: Opens the <i>Create Recurrence</i> dialog box where you can create or edit a scheduled job recurrence pattern.
	Run Now: Runs the scheduled job immediately.
	Delete: Deletes the scheduled job.

4.4 Log Explorer

The Log Explorer allows you to view general information, warnings, or error messages. Right-click any log entry to display a pop-up menu where you can select commands to View Details of a log entry, Delete a log entry, or View XML format of a log entry.

Figure 4-3 *Working with the Log Explorer*

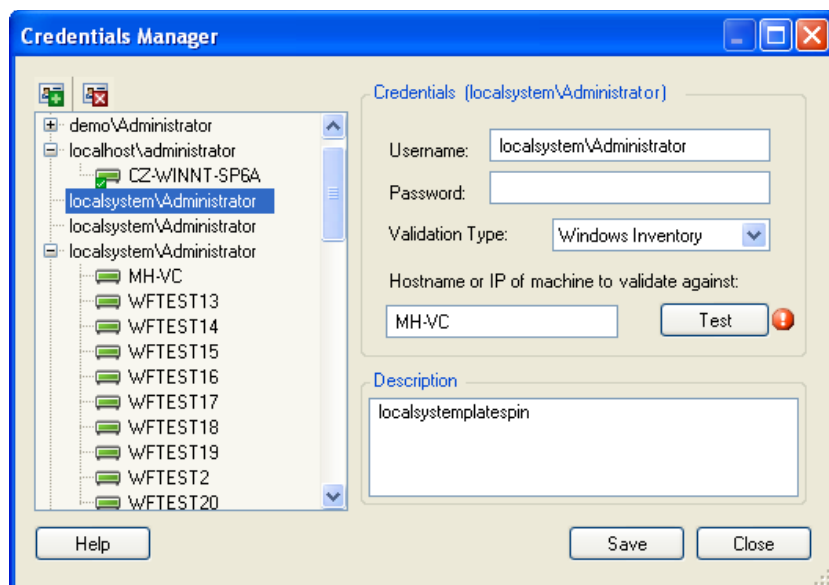


The View Details and Delete commands are also available on the tool bar. To delete multiple logs at once, select the logs using the Shift or Ctrl keys and press Delete. Messages that are displayed here are not related to any particular job or server.


4.5 Credentials Manager

Use the Credentials Manager to create and manage the credentials required to access servers. Create new credentials, delete saved credentials, or edit credentials in the Credentials Manager dialog box.

Figure 4-4 *Credentials Manager dialog box*



To create a new credential:

- 1 Click the *New* button .
- 2 Type in a *Username* and *Password*.

NOTE: If adding a Windows Inventory credential, it is necessary to combine the username with the domain name in the form: <Domain>\<Username> or Username@Domain. For example, platespin\reconuser.

- 3 Select a *Validation Type* from the drop-down listbox. Supported types are: Windows Inventory, Unix Inventory, Virtual Center Inventory, VMware ESX3 Inventory, Novell Network Inventory, Citrix XenServer Inventory, Database or Other. By default, Other is selected.
- 4 (Optional) Type a *Hostname* or *IP Address* to validate the credentials against. Click *Test* to validate the entered credentials.
- 5 (Optional) Enter a *Description* to distinguish the credentials from other credentials with identical user names.
- 6 Click *Save* to store the credentials or click *Close* to exit without saving your changes.

To delete a credential, select it in the left pane and click the *Delete* button .

NOTE: Credentials cannot be deleted if they are assigned to servers. If credentials are assigned to servers, they are a plus sign beside them in the left pane. Click the plus sign (or double-click the credential) to see which machines have those credentials assigned to them.

4.6 PlateSpin Analyzer

The PlateSpin Analyzer quickly analyzes networks and determines whether PlateSpin Recon can successfully inventory, monitor, or convert the specified Windows servers. It performs two different types of tests: Inventory and Conversion.

- ♦ [Section 4.6.1, “Using PlateSpin Analyzer,” on page 71](#)
- ♦ [Section 4.6.2, “Inventory and Monitoring Tests,” on page 72](#)
- ♦ [Section 4.6.3, “Conversion Tests,” on page 73](#)

4.6.1 Using PlateSpin Analyzer

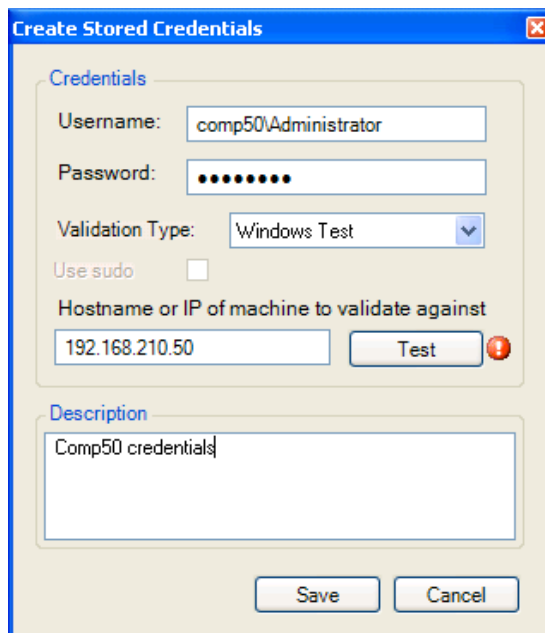
- 1 Click *Tools > Analyze Servers*.

The PlateSpin Analyzer window is displayed.

- 2 Click the *Select Credential* drop-down listbox and select *<New Credential...>*.

The Create Stored Credentials dialog box is displayed.

Figure 4-5 Create Stored Credentials dialog box

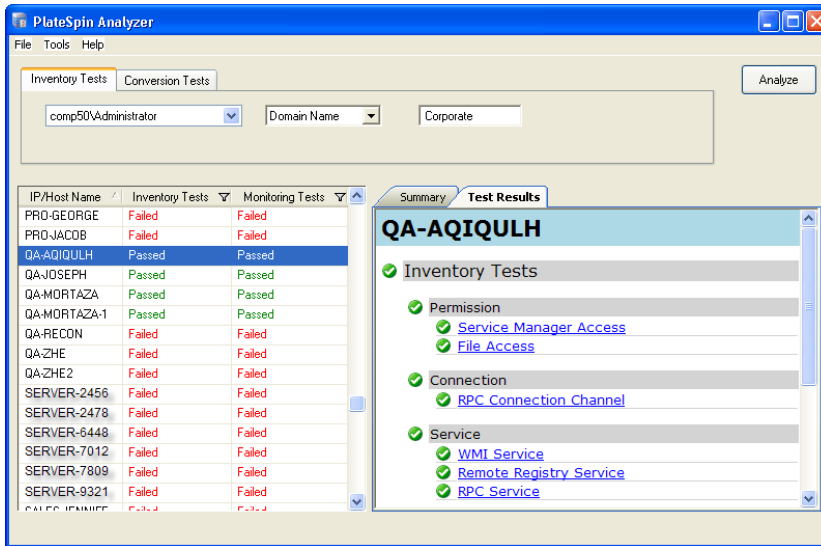


- 3 Type a *Username* and *Password* in the form `domain\username` or `hostname\username`.
- 4 Click the *Validation Type* drop-down box and select either *Windows Test* or *Unix Test*.
PlateSpin Analyzer does not support analysis of Unix/Linux servers, but they can be used for credential validation.
- 5 (Optional) Type a description for the credentials.
- 6 Click *Save*.

NOTE: If *Save* is grayed out, verify that the credentials are entered in the correct format.

- 7 Click the *Domain Name* drop-down listbox and select one of the available options: *Domain Name*, *IP/Host Name*, *IP Range*.
- 8 Type a value in the empty field appropriate for the selection in Step 5. For example, if you selected *Domain Name*, enter your domain name.
- 9 Click *Analyze* and wait for the process to complete.

Figure 4-6 PlateSpin Analyze analysis display



When the analysis is complete, the left pane displays either *Passed* or *Failed*. Click the *Test Results* tab in the right pane to view result details. If any of the tests returned a failed result, the cause must be resolved before PlateSpin Recon can successfully inventory the host. A failed result typically indicates that there are insufficient permissions and/or that the required services are not running.

4.6.2 Inventory and Monitoring Tests

Inventory tests are divided into 3 main groups, each containing a subset of tests:

Table 4-3 Inventory Test groups

Main Test Group	Test Group Subset	Description
Permission	Service Manager Access	Verifies that the user has permission to open the Service Manager.
	File Access	Verifies that the correct access permission is granted on the System32 folder of the machine being inventoried.
Connection	RPC Connection Channel	Determines whether an RPC connection channel can be established by resolving the DNS/NETBIOS name.

Main Test Group	Test Group Subset	Description
Service	WMI Service	Determines whether the WMI (Windows Management Instrumentation) service is enabled and running on the client machine(s) being inventoried.
	Remote Registry Service	Verifies that the Remote Registry Service is enabled on the client machine(s).
	RPC Service	Verifies that the RPC (Remote Procedure Call) Service is enabled on the client machine(s).

Monitoring tests are divided into 2 main groups, each containing a subset of tests:

Table 4-4 *Monitoring test groups*

Service	RPC Service	Verifies that the RPC service is running on the client machine(s).
	DCOM Service	Verifies that the DCOM service is running on the client machine(s).
Port	TCP Port 135	Determines whether TCP port 135 and 445 are available. To obtain more conclusive information, consider running Port Scanner for all TCP ports above 1024.
	TCP Port 445	

4.6.3 Conversion Tests

Conversion Tests are similar to those performed by the PlateSpin Portability Suite PlateSpin Analyzer tool. For more information, refer to Knowledge Base article [Q20478 \(http://support.platespin.com/kb2/article.aspx?id=20478\)](http://support.platespin.com/kb2/article.aspx?id=20478).

Generating and Viewing Charts and Reports

5

When PlateSpin® Recon has collected inventory and monitoring data, the data can be displayed for analysis in charts or customizable reports. The Chart Viewer and the Report Explorer are responsible for displaying collected information in a variety of formats.

- ♦ [Section 5.1, “Chart Viewer,” on page 75](#)
- ♦ [Section 5.2, “Report Explorer,” on page 79](#)
- ♦ [Section 5.3, “Editing Report Templates,” on page 84](#)
- ♦ [Section 5.4, “View Customization,” on page 92](#)
- ♦ [Section 5.5, “Generating Reports,” on page 96](#)
- ♦ [Section 5.6, “Document Viewer,” on page 100](#)

5.1 Chart Viewer

Charts provide a graphical representation of workload and utilization data. To view a chart for a given server, select the server in the Data Center Explorer or click a cell belonging to the server in a report view. You can also select several servers in the Data Center Explorer (hold down Ctrl or Shift while you select the servers) to render charts for multiple servers. In this case, a chart is not automatically rendered, but a *Create Chart* button is displayed in the Chart Viewer. Click *Create Chart* to render the chart.

- ♦ [Section 5.1.1, “Chart Controls and Commands,” on page 76](#)
- ♦ [Section 5.1.2, “Expression Types,” on page 77](#)
- ♦ [Section 5.1.3, “Level of Detail,” on page 78](#)
- ♦ [Section 5.1.4, “Data Type Statistics,” on page 78](#)
- ♦ [Section 5.1.5, “Server Instance,” on page 78](#)
- ♦ [Section 5.1.6, “Date Range Presets,” on page 79](#)
- ♦ [Section 5.1.7, “Forecasting,” on page 79](#)

All times displayed in the Chart Viewer are local times. Times shown in reports are in UTC time.



5.1.1 Chart Controls and Commands

Figure 5-1 Chart Controls



Table 5-1 Chart Viewer Commands

Command	Description
	Print Preview. Preview chart before printing.
	Print. Print the chart.
	Copy To Clipboard. Copy the chart to the clipboard.
	Line. Display straight lines between points of continuous data.
	Curve. Display curved lines between points of continuous data.
	Candlestick. View the variability of sample data. This button appears only for the candle chart.
	Stacked Bar. View cumulative workload data from multiple servers. This is enabled for Expression types other than Utilization.
	Stacked Area. View cumulative workload data from multiple servers. This is enabled for Expression types other than Utilization.
	Legend Box. Give a visual cue and description for each data object. The legend is very useful for quickly glancing through the data on a chart.
	Point Labels. Display labels for the values in a chart.
	Zoom. Switch between normal and magnified view.
	Zoom In/Out. Adjust the granularity of the chart for the collected data. See Section 5.1.3, "Level of Detail," on page 78.
	Navigate Beginning/Back/Forward/End. Scroll the visible data.
	Historical Data Settings. Select the historical data time range upon which to base forecasting.

Command	Description
	Forecast Settings. Select the forecasting time period and algorithm. This button is only enabled for daily or weekly Series charts or for Profile charts.
	Refresh All Data From Server. Refresh the view.

5.1.2 Expression Types

Expressions are individual data series that can be viewed on the currently displayed chart. Select an available Type/Expression combination and the chart is rendered to display the new metrics. The list of available expressions varies depending on the expression type and machine selected.

Table 5-2 *Expression types*

Type	Expressions Available	Description
Workload	Disk (MB/s)	The operating system, applications, and data stack running on a server, with the quantified dimensions of Processor Speed, Memory, Disk Space, Network, and Time.
	Disk (Transfers/sec)	
	Disk Queue Length	
	Disk Space Free (GB)	
	Disk Space Used (GB)	
	Disk Writes (MB/sec)	
	Memory (Pages/sec)	
	Memory Free (MB)	
	Memory Used (MB)	
	Network (MB/sec)	
	Network (Packets/sec)	
	Processor (MHz)	
	Processor (Normalized)	
	Processor Queue Length	
	Total Disk (GB)	
	Total Memory (MB)	
	Total Processor (Count)	
Utilization	Disk IO (%)	Characterizes the fit between workload and resources (% utilization).
	Disk Space Free (%)	
	Disk Space Used (%)	
	Memory Free (%)	
	Memory Used (%)	
	Processor (%)	

Type	Expressions Available	Description
Counter	Counter values are named for the Windows Performance Monitor counters used to retrieve them from Windows systems. The name consists of the Performance Monitor category joined by an underscore to the counter name. For example, Memory/PagesPerSecond becomes Memory_PagesPerSecond.	Primitive measure of performance collected from the operating system.

5.1.3 Level of Detail

From the drop-down list, select from the following levels of granularity (in terms of time) for collected data:

- ♦ **5 Min:** Averages of data collected during five-minute intervals. (Series charts only.)
- ♦ **1 Hour:** Summaries of data collected during one-hour intervals. (Series and Candle charts only.)
- ♦ **1 Day:** Characterizes workloads over a day. This comprises 24 values that correspond to each hour of the day (0 to 23). Each value is calculated by taking the peak value from all values observed for that hour. (All chart types.)
- ♦ **1 Week:** Characterizes workloads over a week. (All chart types.)

5.1.4 Data Type Statistics

The following data type statistics are collected for each hour and are available for selection when a 1 hour or higher zoom level is selected: This option is available only for series charts.

- ♦ **AVG:** Average of all values.
- ♦ **MAX:** Maximum sustained load.
- ♦ **MIN:** Minimum sustained load.
- ♦ **STDEV:** Standard deviation.

NOTE: Standard deviation is not calculated for Virtual Center counters.

- ♦ **Peak AVG:** : (1 Day or 1 Week only) Peak average represents the averages of the values for the highest period for each 12-hour interval or each day in the reporting period.

5.1.5 Server Instance

Click the drop-down list to see all the instances of the currently selected resource (expression). Select each instance to see chart data for that resource. If multiple servers were used to render the chart, all instances of the resources for all servers are displayed. This option is available only for candle charts.



5.1.6 Date Range Presets

Select from the drop-down list *Last 3 days*, *Last 7 days*, *Last Month*, or *All Data*. Select *<Custom>* to open the Advanced Date Filter Settings dialog box for more specific date range options.

5.1.7 Forecasting

Forecasting is available for series and profile charts, for the 1 day and 1 week level of detail. At least three days of monitoring data are necessary to use forecasting in charts.

To add forecasting to a chart view:

- 1 Display a *Profile* chart.
- 2 Click the *Date Range* drop-down list and select a time period to base the forecast on .
- 3 (Optional) Click *Historical Data Settings*  for advanced date filter settings.
The Advanced Date Filter Settings dialog box is displayed. Choose any desired settings and click *OK*.
- 4 Click the *Forecast Period* drop-down list and select a time period for the forecast .
- 5 (Optional) Click *Forecast Settings*  for advanced custom forecast settings.
The Advanced Forecast Settings dialog box is displayed. Choose any desired settings and click *OK*.
- 6 Click the *Forecast > Trend* drop-down list and select a forecast trend or regression method.

5.2 Report Explorer

The Report Explorer is used to organize and generate reports.

- ♦ [Section 5.2.1, “Report Explorer Template Nodes,” on page 79](#)
- ♦ [Section 5.2.2, “Commands,” on page 82](#)
- ♦ [Section 5.2.3, “Analysis Reports,” on page 83](#)
- ♦ [Section 5.2.4, “Working with Report Templates,” on page 83](#)

5.2.1 Report Explorer Template Nodes

The Report Explorer hierarchy consists of template nodes that have various types of report templates.

Table 5-3 *Report Explorer Template Nodes*

Template Node	Report Templates	Node Description
Analysis	Application - Exchange Workloads	For reporting dynamic data such as hardware utilization and workload sizes.
	Application - SQL Server Workloads	
	Disk - Average 30 Day Forecast	
	Disk - Peak	
	Hardware Utilization - Average 30 Day Forecast	
	Hardware Utilization - Peak	
	Memory - Average 30 Day Forecast	
	Memory - Overutilized	
	Memory - Peak	
	Memory - Underutilized	
	Network -Average 30 Day Forecast	
	Network - Peak	
	Processor - Average 30 Day Forecast	
	Processor - Overutilized	
	Processor - Peak	
	Processor - Underutilized	
	Resource Utilization Trend	
Chargeback	Cost Allocation	For reporting the transfer of charges incurred by the IT department back to the business units that are using services. See Section 7.4, "Chargeback Reporting," on page 146.
Consolidation	Project	For reporting on projects, scenarios, and servers in the Consolidation Planning Module. See Section 6.5, "Viewing Consolidation Reports," on page 132 for information on how to use templates within this node.
	Scenario Summary	
	Workload Assignment	

Template Node	Report Templates	Node Description
Consolidation Project Samples	Good Consolidation Candidates	Commonly used sample consolidation reports.
	Inventory Summary	
	Poor Consolidation Candidates	
	Utilization Summary	
	Workload List with Conditional Formatting	
Diagnostics	Error Report	For diagnosing errors that cause monitoring data to be missing from a chart.
	Missing Data - Detailed	
	Missing Data - Summary	
Disk	Disk Inventory	For reporting the size and type of physical disks on machines.
Inventory	Hardware - Obsolete	For reporting static data such as operating system, processor speed, memory, disk space, network, power, and cooling.
	Hardware Inventory	
	Power and Cooling	
	Virtual Machine Server Candidates	
	Virtual Machine State	
	VM Files Report	
	VM Snapshots Report	
Reclamation	Disk Reclamation	For reporting on elements that can be reclaimed to ensure efficient use of resources.
	Memory Reclamation	
	Powered On VMs Reclamation	
	Volume Reclamation	
Software	Application > Application Inventory	For reporting installed applications, services, and running processes.
	Application Complete > Application Complete Inventory	
	Processes > Process Inventory	
	Service > Service Inventory	
	Service Complete > Service Complete Inventory	







Template Node	Report Templates	Node Description
Virtual Infrastructure	Cluster Configuration	For reporting on virtual machines and virtual machine hosts.
	Resource Pool Configuration	
	VI Configuration	
	VirtualCenter Report	
	VM Cluster	
	VM Details Report	
	VM Resource Pool	
	VM Virtual Infrastructure	
Volume	Volume Inventory	For reporting the space and file system associated with each logical volume on machines.

As shown in [Table 5-3 on page 80](#), within each report node are one or more templates that you can use as is or that you can customize with sorting, filtering, and view selection options. You can add additional Analysis, Chargeback, Disk, Inventory, Software, and Volume report templates.

5.2.2 Commands

The following Report Explorer commands are accessible by right-clicking a report template or a generated report:

Table 5-4 Report Explorer Commands

Commands	Tooltip/Description
	Create Report Template. Create a new report template based on the report node selected.
	Copy Report Template. Copy a report template.
	Edit Report Template. Edit a report by setting group, views, filters, summaries, and forecasts for a report template or a report.
	Generate Report. Generate a report based on the report template selected.
	View Report. View an existing (previously generated) report.
	Send Report. Send a generated report via e-mail or to an FTP server or network location.
	Delete. Delete the selected report or template.

5.2.3 Analysis Reports

Analysis reports use dynamic data such as hardware utilization and workload sizes. To view workloads or utilization of servers that are being monitored, generate an Analysis report. The report is displayed with historical data, which is useful for getting a single value (maximum, minimum, peak, average) for expressions over a specified period of time. Use this data to identify over-utilized or under-utilized server resources.

NOTE: Data must be collected for at least one hour before Analysis reports are generated. If there is insufficient data, the report is displayed without dynamic data.

To get detailed historical data for each data point, click the report cell to view the related detailed chart.

Machines without monitoring data are included in Analysis reports by default to allow complete analysis of all inventoried machines.

Figure 5-2 Sample Analysis Report

Hardware Utilization - Peak							
Machine	Operating System	Memory			Processor		
Host	Type	Size (MB)	Used (%)		Total Speed (MHz)	%	
			PEAK	Hour		PEAK	Hour
DEV-YONINAJ	Windows XP	2048.0	91.3	14	3990.0	55.2	13
DEV-FREETS	Windows 2003	2048.0	36.8	16	5984.0	2.8	16
DEV-NADERS	Windows 2003	2048.0	43.4	18	3990.0	3.8	19
DEV-CYRILM	Windows 2003	2040.0	15.2	14	3740.0	3.4	13
DEV-TONVYL	Windows 2003	3584.0	61.8	16	6384.0	25.3	18
DEV-MICHAELB	Windows XP	2048.0	94.1	18	6384.0	12.2	13
DEV-MIRAG	Windows 2003	2048.0	93.6	19	5984.0	27.8	19
DEV-EDWARDC	Windows 2003	2048.0	37.4	15	3990.0	1.6	15
DEV-TONYP	Windows 2003	2048.0	85.2	21	3990.0	17.1	18
DEV-MIRCEAA	Windows XP	2048.0	56.5	20	3994.0	26.5	16
DEV-ARIO	Windows 2003 (64bit)	4032.0	71.6	15	5320.0	31.3	13
DEV-WILSONT	Windows 2003	2048.0	90.6	20	3740.0	39.9	10
DEV-GEORGED	Windows 2003	3072.0	73.2	19	3990.0	24.1	13

5.2.4 Working with Report Templates

- “Creating a Report Template” on page 83
- “Copying a Report Template” on page 84
- “Deleting a Report Template” on page 84

Creating a Report Template

- 1 Right-click the report type node in the Report Explorer for the type of report that you want to create, then select *Create Report Template* from the menu.

The Create Report Template dialog box is displayed.

- 2 Define the template settings, using the procedures and guidelines for editing an existing template. See [Section 5.3, “Editing Report Templates,” on page 84](#).
- 3 Click *OK*.

The new template is displayed under the report type node in the Report Explorer.

Copying a Report Template

- 1 Right-click the report template you want to copy in the Report Explorer and select *Copy Report Template* from the menu.

The Edit Report Template dialog box is displayed with (COPY) added to the template name.

- 2 (Optional) Edit the template name.
- 3 Click *OK*.

The new copy of the template is displayed in the Report Explorer.

Deleting a Report Template

- 1 Right-click the report template you want to delete in the Report Explorer and select *Delete* from the menu.

The Confirm Delete dialog box is displayed.

- 2 Click *OK*.

The report template is removed from the Report Explorer.

5.3 Editing Report Templates

PlateSpin Recon includes a library of standard, predefined templates available for generating reports. Templates specify the content to include in a report and the form in which the content is presented.

- ♦ [Section 5.3.1, “Report Template Option Tabs,” on page 85](#)
- ♦ [Section 5.3.2, “Chargeback Report Templates,” on page 89](#)
- ♦ [Section 5.3.3, “Consolidation Report Templates,” on page 89](#)
- ♦ [Section 5.3.4, “Consolidation Project Samples Report Templates,” on page 90](#)
- ♦ [Section 5.3.5, “Diagnostics Report Templates,” on page 90](#)
- ♦ [Section 5.3.6, “Disk Report Templates,” on page 90](#)
- ♦ [Section 5.3.7, “Inventory Report Templates,” on page 91](#)
- ♦ [Section 5.3.8, “Reclamation Report Templates,” on page 91](#)
- ♦ [Section 5.3.9, “Software Report Templates,” on page 91](#)
- ♦ [Section 5.3.10, “Virtual Infrastructure Report Templates,” on page 91](#)
- ♦ [Section 5.3.11, “Volume Report Templates,” on page 92](#)

PlateSpin Recon reporting capabilities allow for the customization of most report templates to meet specific business requirements. You can configure the content and summaries included in most reports, as well as the time frame, report presentation, and a number of other elements. Users can save modified templates and generate reports from them.

For all reports except Consolidation reports, you can copy the template and then edit each copy. This preserves the original. Use the modified templates to generate alternate reports.

To edit a report template, right-click the template and select *Edit Report Template* from the menu. The Edit Report Template dialog box is displayed.

5.3.1 Report Template Option Tabs

The options available for customizing report templates vary depending on the type of report template selected.

- ♦ “Group Tab” on page 85
- ♦ “Views Tab” on page 85
- ♦ “Field Filter Tab” on page 85
- ♦ “Row Filter Tab” on page 87
- ♦ “Columns Tab” on page 87
- ♦ “Time Period Tab” on page 87
- ♦ “Options Tab” on page 88
- ♦ “Date Filter Tab” on page 88
- ♦ “Summaries Tab” on page 88
- ♦ “Forecast Tab” on page 89

Group Tab

Allows you to select filters or groups of servers to include in the report. See [Section 3.4, “Working with Groups,” on page 29](#) for more information.

Views Tab

This tab displays the views available for the report, including default and previously created views. Use the arrow buttons to reorder the views. Use the check boxes in the *Visible* column to determine which views to display in the report.

To add another view to a report, click *Add*. To edit an existing view, select it and click *Edit*. For more information, see [Section 5.4.1, “Creating and Editing Report Views,” on page 92](#). To delete a view, select it and click *Delete*.

Field Filter Tab

You can filter servers based on the constraints you set and use expressions to specify how the constraints should be used and interpreted to narrow or expand the report contents.

Figure 5-3 Field Filter Tab

To add a query statement:

- 1 Enter query constraints in the fields and drop-down lists at the top of the Field Filter page.
There might not be a units field or drop-down list, depending on the resource selected in the first drop-down list.
- 2 Click *Add*.
The query statement is added to the list of queries.

To edit a query statement:

- 1 Select the query statement.
- 2 Click *Edit*.
The query settings are copied to the fields at the top of the page .
- 3 Edit the settings as desired.
- 4 Click *Add*.
The selected query statement is updated with the new settings.

To remove a query statement:

- 1 Select the query statement you want to delete.
- 2 Click *Remove*.
The selected query statement is deleted.

The filter expression determines how the query statements are evaluated.

Filters affect which servers are made available for a report views.

Table 5-5 *Filter Expressions*

Filter Expression	Description
AND	All query statements must be true for a server to be added to the gathered pool of servers. For example, if there are several query statements, the first one is evaluated to gather some servers. Then the second query statement is evaluated against the gathered servers, removing any from the pool that don't agree with the new query statement. This continues through all the query statements.
OR	All query statements are evaluated individually. For example, if there are several query statements, the first one is evaluated to gather some servers. Then the second query statement is evaluated adding any new servers not gathered by the first statement to the pool of gathered servers. This continues through all the query statements.
Custom	A Custom expression allows you to specify how to evaluate the query statements. A mixture of OR and AND criteria can be used. You can use parentheses to control the evaluation. Use the query statement ID numbers in your expression. If Custom is selected, the accompanying field cannot be empty.

The *Filter out empty values* check box excludes machines from the report if they have fields with null values.

Row Filter Tab

Used to determine which rows are included or excluded from the report. This is different from *Field Filter* tabs, which control which machines are included or excluded from the report, but the guidelines for managing Field Filter criteria can also be applied here See [“Field Filter Tab” on page 85](#).

This option is available only for the Application Complete and Software Complete reports.

Columns Tab

Allows you to customize the columns displayed in the reports generated, based on these templates. See [Step 1 on page 101](#) for more information on the settings in this page.

This option is available only for the software report templates.

Time Period Tab

Allows you to specify the time period to use for these reports. The settings on this page are the same as the top part of the Date Filter page . See [“Date Filter Tab” on page 88](#) and [Section 5.3.5, “Diagnostics Report Templates,” on page 90](#).

This option is available only for the diagnostic report templates.

Options Tab

Allows you to specify the counter setting report mode for these reports. See [Section 5.3.5, “Diagnostics Report Templates,”](#) on page 90.

This option is available only for the missing data diagnostic report templates.

Date Filter Tab

Use this tab to set the date and time range for calculating summary data values in a report. For example, select the Last 1 Month(s) collection range, Monday to Friday, and 9:00 AM to 5:00 PM to view workload data summarized over business hours for all servers that are monitored:

Figure 5-4 Date Filter Tab

The screenshot shows the 'Edit Report Template' dialog box with the 'Date Filter' tab selected. The 'Name' field contains 'Hardware Utilization - Prime Time Average'. The 'Date Range' section has 'Last' selected with a value of '1' and 'Month(s)'. The 'Business Days' section has checkboxes for Monday, Tuesday, Wednesday, Thursday, and Friday, all of which are checked. The 'Business Hours' section has a checkbox 'Use the following hour range:' which is checked, with 'From' set to '09:00 AM' and 'To' set to '05:00 PM'. At the bottom are 'Help', 'OK', and 'Cancel' buttons.

Summaries Tab

Summarizes average (AVG), maximum (MAX), minimum (MIN) and standard deviation (STDEV) data types. AVG data can be calculated by Average, Peak Hour, or Total of the data collected.

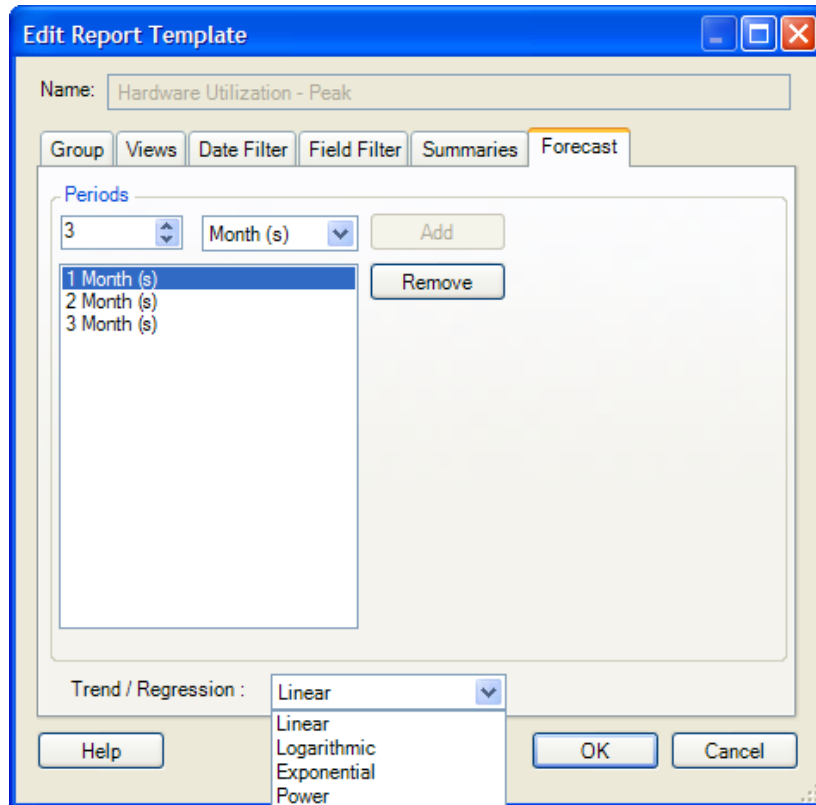
Calculation types:

- ♦ **Average:** Average of all hourly values across business days and business hours.
- ♦ **Peak Hour:** A calculation based on the highest hourly value.
- ♦ **Total:** Sums all the data points in the time frame. For example, Disk IOs per second, for a 1-hour period equates to 10 disk IOs per second X 60 seconds X 60 minutes = 36000. Select the *Summary Rows* column in the report to display the columns containing calculated data.

Forecast Tab

Projects workloads and utilization into the future based on summary data for a historical period. This lets data center architects anticipate future needs and plan accordingly.

Figure 5-5 *Forecast Tab*



5.3.2 Chargeback Report Templates

Edit Chargeback report templates in the same way as Analysis report templates. See [Section 5.3.1, “Report Template Option Tabs,”](#) on page 85 for information about the *Group*, *Views*, and *Field Filter* tabs.

Chargeback reports support only table views. See [“Creating and Editing a Table View”](#) on page 93 for information about table views. See [Section 5.2.4, “Working with Report Templates,”](#) on page 83 for information about the *Date Filter*, *Summaries* and *Forecast* tabs.

5.3.3 Consolidation Report Templates

Project, Scenario Summary, and Workload Assignment reports are based on editable templates. Unlike Inventory and Analysis report templates, there is only one version of a given Consolidation report template available at a time. Edit the template as desired before viewing the report.

- 1 In the Report Explorer, right-click the desired report template, then click *Edit Report Template*. The appropriate Default Settings dialog box is displayed.
- 2 Select the required report columns, data sorting options, and summary rows, then click *Save*.

For more detailed information on the other configuration options of the consolidation report template, see [Section 5.6.2, “Customizing Reports,” on page 101](#).

5.3.4 Consolidation Project Samples Report Templates

Consolidation Project Samples reports are based on editable templates. The templates can be copied or edited. Edit the template as desired before viewing the report.

- 1 In the Report Explorer, right-click the desired report template, then click *Edit Report Template*. The appropriate Default Settings dialog box is displayed.
- 2 Select the required report columns, data sorting options, and summary rows, then click *Save*. For information on the option tabs, see [Section 5.3.1, “Report Template Option Tabs,” on page 85](#).

For more detailed information on the other configuration options of the Consolidation Project Samples report templates, see [Section 5.6.2, “Customizing Reports,” on page 101](#).

5.3.5 Diagnostics Report Templates

Edit Diagnostics report templates by changing options in the Edit Report Template dialog box. To access this dialog box, right-click a Diagnostics report template in the Report Explorer, and on the menu, click *Edit Report Template*.

The available dialog box tabs depend on the *Diagnostics* report template selected:

- ♦ **Group:** See [“Group Tab” on page 85](#).
- ♦ **Views:** See [Section 5.4.1, “Creating and Editing Report Views,” on page 92](#).
- ♦ **Time Period:** Select a time period upon which to base the report. Any intervals during that time period that are missing data are included in the report.
- ♦ **Options:** For Missing Data reports only. Select a Report Mode to determine which intervals to include in the report:
 - ♦ **Online:** Includes intervals during which all counters are missing data.
 - ♦ **CPM:** Includes intervals during which CPM counters are missing data.
 - ♦ **Full:** Includes intervals during which any counter is missing data.

5.3.6 Disk Report Templates

Disk reports are based on editable templates. The templates can be copied or edited. Edit the template as desired before viewing the report.

- 1 In the Report Explorer, right-click the desired report template, then click *Edit Report Template*. The appropriate Default Settings dialog box is displayed.
- 2 Select the required report columns, data sorting options, and summary rows, then click *Save*. For more information on the option tabs, see [Section 5.3.1, “Report Template Option Tabs,” on page 85](#).

For more detailed information on the other configuration options of the Disk report templates, see [Section 5.6.2, “Customizing Reports,” on page 101](#).

5.3.7 Inventory Report Templates

Inventory reports are based on editable templates. The templates can be copied or edited. Edit the template as desired before viewing the report.

- 1 In the Report Explorer, right-click the desired report template, then click *Edit Report Template*. The appropriate Default Settings dialog box is displayed.
- 2 Select the required report columns, data sorting options, and summary rows, then click *Save*. For more information on the option tabs, see [Section 5.3.1, “Report Template Option Tabs,” on page 85](#).

For more detailed information on the other configuration options of the Inventory report templates, see [Section 5.6.2, “Customizing Reports,” on page 101](#).

5.3.8 Reclamation Report Templates

Reclamation reports are based on editable templates. The templates can be copied or edited. Edit the template as desired before viewing the report.

- 1 In the Report Explorer, right-click the desired report template, then click *Edit Report Template*. The appropriate Default Settings dialog box is displayed.
- 2 Select the required report columns, data sorting options, and summary rows, then click *Save*.

For more detailed information on the other configuration options of the Reclamation report templates, see [Section 5.6.2, “Customizing Reports,” on page 101](#).

5.3.9 Software Report Templates

Software reports are based on editable templates. The templates can be copied or edited. Edit the template as desired before viewing the report.

- 1 In the Report Explorer, right-click the desired report template, then click *Edit Report Template*. The appropriate Default Settings dialog box is displayed.
- 2 Select the required report columns, data sorting options, and summary rows, then click *Save*. For more information on the option tabs, see [Section 5.3.1, “Report Template Option Tabs,” on page 85](#).

For more detailed information on the other configuration options of the Software report templates, see [Section 5.6.2, “Customizing Reports,” on page 101](#).

5.3.10 Virtual Infrastructure Report Templates

Virtual Infrastructure reports are based on editable templates. The templates can be copied or edited. Edit the template as desired before viewing the report.

- 1 In the Report Explorer, right-click the desired report template, then click *Edit Report Template*. The appropriate Default Settings dialog box is displayed.
- 2 Select the required report columns, data sorting options, and summary rows, then click *Save*. For more information on the option tabs, see [Section 5.3.1, “Report Template Option Tabs,” on page 85](#).

For more detailed information on the other configuration options of the Virtual Infrastructure report templates, see [Section 5.6.2, “Customizing Reports,” on page 101](#).

5.3.11 Volume Report Templates

Volume reports are based on editable templates. The templates can be copied or edited. Edit the template as desired before viewing the report.

- 1 In the Report Explorer, right-click the desired report template, then click *Edit Report Template*. The appropriate Default Settings dialog box is displayed.
- 2 Select the required report columns, data sorting options, and summary rows, then click *Save*. For more information on the option tabs, see [Section 5.3.1, “Report Template Option Tabs,” on page 85](#).

For more detailed information on the other configuration options of the Volume report templates, see [Section 5.6.2, “Customizing Reports,” on page 101](#).

5.4 View Customization

This section discusses the customization of report views, specifically table and chart views.

- ♦ [Section 5.4.1, “Creating and Editing Report Views,” on page 92](#)
- ♦ [Section 5.4.2, “Histogram Charts,” on page 94](#)
- ♦ [Section 5.4.3, “Scatter Charts,” on page 95](#)
- ♦ [Section 5.4.4, “Usage Charts,” on page 96](#)

5.4.1 Creating and Editing Report Views

- ♦ [“Adding a View to a Report” on page 92](#)
- ♦ [“Editing a Report View” on page 93](#)
- ♦ [“Removing a View from a Report” on page 93](#)
- ♦ [“Creating and Editing a Table View” on page 93](#)
- ♦ [“Creating and Editing a Chart View” on page 93](#)

Adding a View to a Report

- 1 From the Report Explorer, right-click the report template you want to customize and select *Edit Report Template*.
The Edit Report Template dialog box is displayed.
- 2 Click the *Views* tab.
The Views page is displayed, showing any default or previously created views.
- 3 Click *Add*.
The Create View (Step 1: General Settings) dialog box is displayed.
- 4 Type a unique *Name* for the new view.
- 5 For *Chart* views, proceed to [“Creating and Editing a Chart View” on page 93](#). For *Table* views, proceed to [“Creating and Editing a Table View” on page 93](#).

Editing a Report View

- 1 Select a view on the *Views* page of the Edit Report Template dialog box.
- 2 Click *Edit*.
The Create View (Step 2: Customization) page is displayed for the view type selected.
- 3 For *Chart* views, proceed to [“Creating and Editing a Chart View” on page 93](#). For *Table* views, proceed to [“Creating and Editing a Table View” on page 93](#).

Removing a View from a Report

- 1 Select a view on the *Views* page of the Edit Report Template dialog box.
- 2 To temporarily remove the view, deselect the *Visible* check box. To permanently remove the view, click *Delete*.
- 3 Click *OK*.

Creating and Editing a Table View

- 1 Select *Table* in the *Type* area and click *Next*.
The Step 2: Customization page is displayed.
- 2 On the *General* tab, select the column titles to display in the report.
See [Section 5.6.2, “Customizing Reports,” on page 101](#) for information on arranging columns, conditional formatting, column sorting, and summary rows.
- 3 Use the *Forecast* tab to create forecast columns in the report view.
Create forecast columns and display: Clear the check box if forecast columns are not required. When selected, this option enables the following column options:
 - ♦ **Original Values:** Displays the original values.
 - ♦ **Change (%):** Displays the percentage of change from the original values.
 - ♦ **Change (Units):** Displays the change from original values in units.Forecasting is only supported in table views.
- 4 Use the *Groups* tab to group rows in the table view. For information on report groups, see [“Adding or Modifying Groups” on page 103](#).
- 5 Use the *Matrix* tab to add Matrix columns to the table view. For information on report matrixes, see [“Adding or Modifying Matrixes” on page 103](#).
- 6 Click *Finish* to add the view to the list of available views in the Edit Report Template dialog box.

Creating and Editing a Chart View

- 1 Select *Chart* in the *Type* area.
- 2 Select a chart type from the drop-down list (histogram, scatter chart, or usage chart. For analysis reports only).
- 3 Select an expression from the drop-down list. The expressions listed vary depending on the chart type selected.

- 4 Click *Next* to display the view. For the histogram chart type, the Settings dialog box is displayed. Use it to set various options as described in [Section 5.4.2, “Histogram Charts,”](#) on [page 94](#).
- 5 Click *OK*.
The Create View dialog box (Step 2: Customization page) is displayed, showing the view.
 - ♦ For a *Scatter* chart, the Chart View Settings dialog box is also displayed. Use it to set various data, summary, and sort options as described in [Section 5.4.3, “Scatter Charts,”](#) on [page 95](#).
 - ♦ For an Analysis report, the *Usage* chart type is also available. Select *Usage* and click *Next*. The Usage View Settings dialog box is displayed. Use it to set various options as described in [Section 5.4.4, “Usage Charts,”](#) on [page 96](#).
- 6 Click *Customize* to return to the chart-appropriate Settings dialog box and modify the specified options.
- 7 Use the various buttons at the top of the chart in the Create View dialog box to print or modify the chart.
- 8 Click *Finish*.


The view is displayed in the list of available views in the Edit Report Template dialog box.

5.4.2 Histogram Charts

In the Create View (Step 2: Customization) dialog box, click *Customize* to modify the contents of the histogram.

The Settings dialog box is displayed.

Figure 5-6 Settings Dialog Box



The default setting for the histogram includes five bins with a Step size of 1.

Auto

Select *Auto* to specify a different number of Bins or Step sizes to be used in the histogram. Click *From palette* to automatically generate colors for the histogram bars from the palette. Click *Default color* to choose one default color for all histogram bars.

Customize

Select *Customize* to modify the upper or lower boundaries of any bin needing modification. Optionally, specify custom labels and colors for the corresponding histogram bars.

Click the *Insert* button to insert a new bin below the currently highlighted bin. The appropriate *Start*, *End*, and *Label* columns are automatically adjusted to accommodate the new bin, which can be customized.

Click the *Remove* button to remove the currently highlighted bin. The appropriate *Start*, *End*, and *Label* columns are automatically adjusted.


Select any number of consecutive or non-consecutive bins and click the *Group* button to group them together. Enter the group's name in the *Group* column for any one of the bins. The name is automatically assigned to all members of the group.

Click the check box to display the grouped bins as a single bin in the chart.

5.4.3 Scatter Charts

In the Create View (Step 2: Customization) dialog box, click the *Customize* button to modify the contents of the scatter chart. The Chart View Settings dialog box is displayed.

Figure 5-7 Chart View Settings



In the *Data* area, choose how much data to include in the chart. Select *All* to include all data, or select *Top* or *Bottom* to limit the chart to the highest or lowest values of data. For *Top* or *Bottom*, specify the number of values, up to 25, to include.

In the *Summary* area, select *Mean* and/or *Median* to include those lines in the chart. These values are based on all data, not just displayed values.

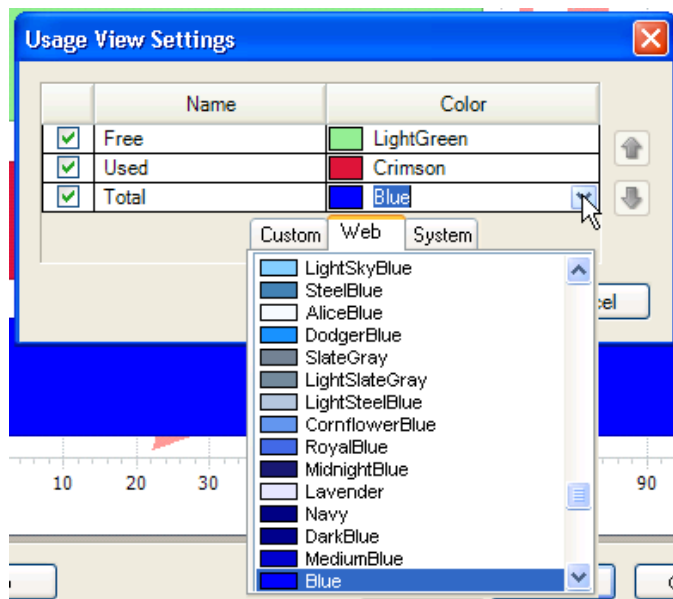
In the *Sort* area, choose to sort the chart in ascending or descending order according to data quantity, or alphabetically by hostname.

Click *OK* to display the chart.

5.4.4 Usage Charts

Usage charts are unique to Analysis report templates. In the Create View (Step 2: Customization) dialog box, click *Customize* to modify the contents of the usage chart. The Usage View Settings dialog box is displayed.


Figure 5-8 Usage View Settings



Use the check boxes to determine which bars to include in the chart. If desired, use the drop-down lists in the *Color* section to assign colors to the bars in the chart.

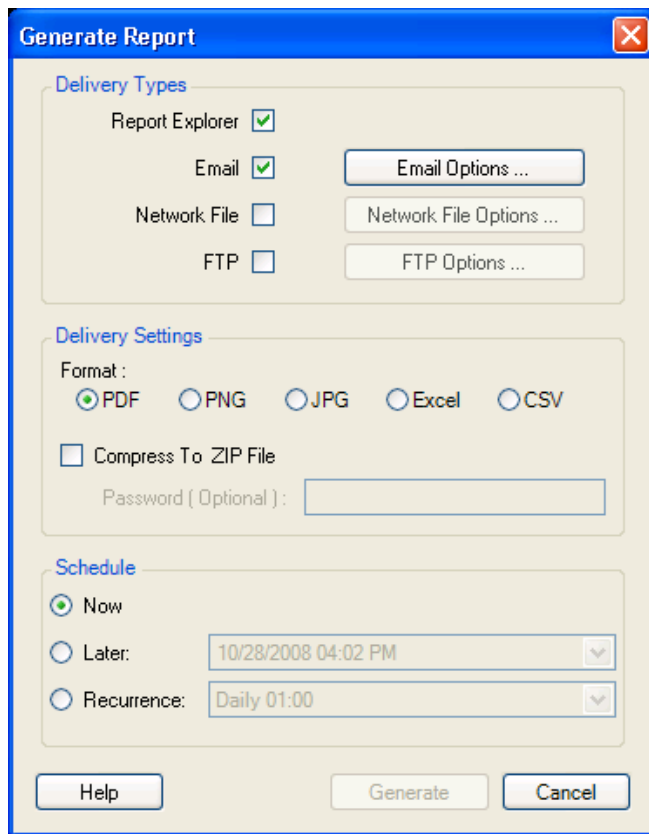
To move a bar to the left or to the right in the report, select it and use the up- or down-arrow. Click *OK* to display the chart.

5.5 Generating Reports

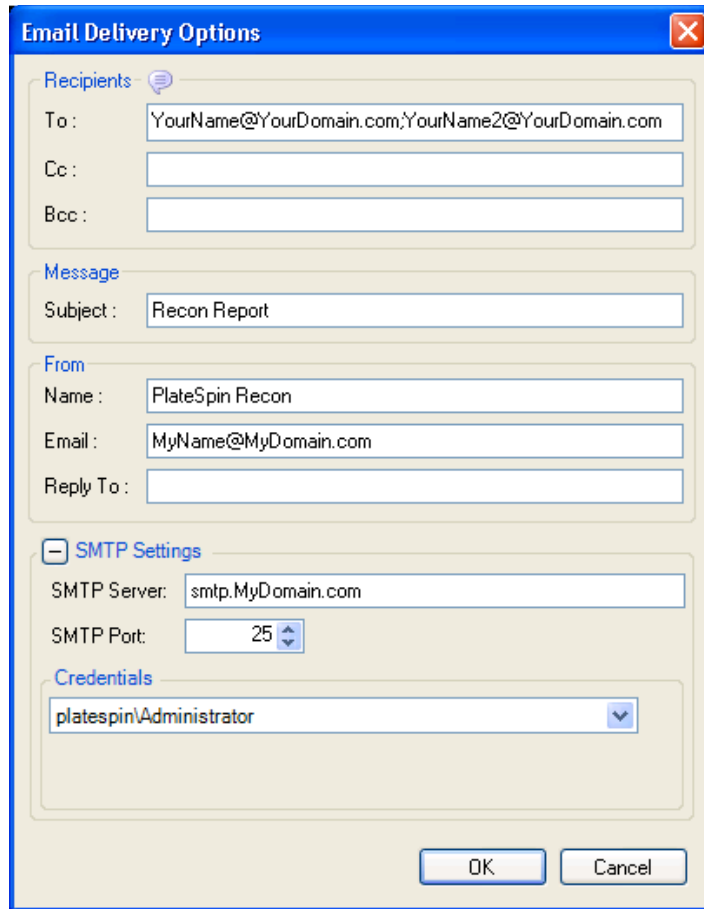
- 1 Select the report template and click *Generate Report* .

The Generate Report dialog box is displayed.

Figure 5-9 *Generate Report Dialog Box*

The image shows a 'Generate Report' dialog box with a blue title bar and a close button (X) in the top right corner. The dialog is divided into three main sections: 'Delivery Types', 'Delivery Settings', and 'Schedule'. In the 'Delivery Types' section, 'Report Explorer' and 'Email' are checked, while 'Network File' and 'FTP' are unchecked. To the right of these checkboxes are buttons for 'Email Options ...', 'Network File Options ...', and 'FTP Options ...'. The 'Delivery Settings' section has a 'Format' section with radio buttons for 'PDF' (selected), 'PNG', 'JPG', 'Excel', and 'CSV'. Below this is a 'Compress To ZIP File' checkbox (unchecked) and a 'Password (Optional)' text field. The 'Schedule' section has radio buttons for 'Now' (selected), 'Later:', and 'Recurrence:'. The 'Later:' option has a date/time dropdown showing '10/28/2008 04:02 PM'. The 'Recurrence:' option has a dropdown showing 'Daily 01:00'. At the bottom of the dialog are three buttons: 'Help', 'Generate', and 'Cancel'.

- 2 Select a delivery types for the report. By default, the report is delivered to the Report Explorer. Other delivery options include:
 - ♦ Sending it by e-mail
 - ♦ Saving it to a network file
 - ♦ Uploading it to an FTP server
- 3 Configure delivery options:
 - ♦ For the Email delivery option, click *Email Options*. The Email Delivery Options dialog box is displayed.
 - ♦ Provide the requested information. Only one recipient is required in the *To* field, although if you enter more than one, you need to separate the addresses with a semicolon. The *CC* and *BCC* fields can be left blank. Any default values supplied for the *Subject* and *Name* can be changed. The *From* Fddress is mandatory, but the *Reply To* address can be left blank. The *SMTP Server* and *SMTP Port* are mandatory. *Credentials* are only optional if they are not required for your SMTP server, which is unlikely.



The dialog box is titled "Email Delivery Options" and contains several sections for configuring email delivery. The "Recipients" section has fields for "To:", "Cc:", and "Bcc:", with the "To:" field containing "YourName@YourDomain.com;YourName2@YourDomain.com". The "Message" section has a "Subject:" field containing "Recon Report". The "From" section has fields for "Name:" (PlateSpin Recon), "Email:" (MyName@MyDomain.com), and "Reply To:". The "SMTP Settings" section has a minus icon, "SMTP Server:" (smtp.MyDomain.com), and "SMTP Port:" (25). The "Credentials" section has a dropdown menu showing "platespin\Administrator". At the bottom are "OK" and "Cancel" buttons.

Email Delivery Options

Recipients

To : YourName@YourDomain.com;YourName2@YourDomain.com

Cc :

Bcc :

Message

Subject : Recon Report

From

Name : PlateSpin Recon

Email : MyName@MyDomain.com

Reply To :

SMTP Settings

SMTP Server: smtp.MyDomain.com

SMTP Port: 25

Credentials

platespin\Administrator

OK Cancel

- ♦ For the *FTP* and *Network File* delivery types, use the adjacent option buttons to open dialog boxes and specify an FTP path or UNC/local folder path. Choose to have later recurrences overwrite the report saved at that location, or append the date and time to the specified filename so that previous reports are preserved.

FTP Delivery Options

File Options

Path : Z:\FTP\

When transferring files :

☒ Always append date and time to filename

☐ Always overwrite existing files

FTP Server

Host : ftp.acme.com

Port : 21

☐ Use Passive Mode

Protocol

☒ FTP ☐ FTPS ☐ SFTP

Credentials

MB-WIN2K3SP1\Administrator

OK Cancel

Network File Delivery Options

UNC or Local Folder Path : Z:\PowerRecon Reports\

Credentials

MB-WIN2K3SP1\Administrator

When transferring files :

☒ Always append date and time to filename

☐ Always overwrite existing files

OK Cancel

- ♦ For delivery options other than Report Explorer, the default format for the report is *PDF*. Additionally, there are several other available delivery formats.
- 4 If desired, select the option to compress the report in a ZIP file. Optionally, the file can be password-protected.
- 5 You can also set options in the *Schedule* area to specify whether to generate the report immediately, at a later time, or at a recurring interval. For information about scheduling options, see [Section 8.2, “Scheduling and Recurrence,” on page 150](#).
- 6 Click *Generate*.

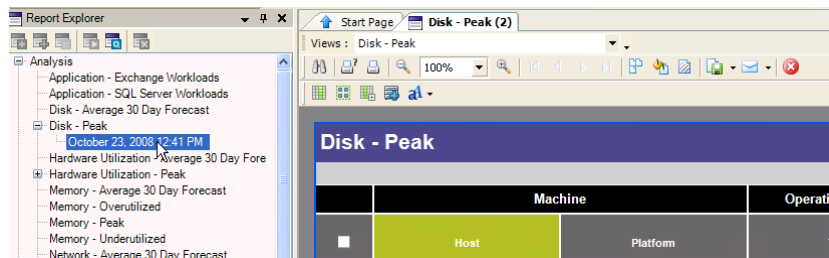
When the report is generated, it is delivered using the delivery types specified. If this includes Report Explorer, the report is cached for quick viewing at any time. An item is added to the Report Explorer beneath the report template used to generate the report, named with the date and time the report was generated.

To view a report, double-click the generated instance. The report is displayed in the Document Viewer. For information, see [Section 5.5, “Generating Reports,” on page 96](#). To delete a report, right-click the report and select *Delete Report* on the menu.

5.6 Document Viewer

To display a report in the Document Viewer, generate a report from a report template and then double-click the instance of the report or double-click a report instance you previously generated. See [Section 5.5, “Generating Reports,” on page 96](#).

Figure 5-10 Report Displayed in the Document Viewer










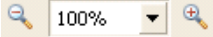






A new tabbed page is added to the Document Viewer, showing the report through its first available view. To switch to a different view, click the *Views* drop-down list and select a view. For more information on views, see [Section 5.4.1, “Creating and Editing Report Views,” on page 92](#).

- ♦ [Section 5.6.1, “Commands,” on page 100](#)
- ♦ [Section 5.6.2, “Customizing Reports,” on page 101](#)
- ♦ [Section 5.6.3, “Viewing Reports,” on page 104](#)
- ♦ [Section 5.6.4, “Printing and Exporting Reports,” on page 105](#)
- ♦ [Section 5.6.5, “Copying Servers from a Report to a Group,” on page 106](#)

5.6.1 Commands

Table 5-6 Document Viewer Commands

Commands	Tooltip/Description
	Columns. Select columns content for viewing.
	Groups. Add, or modify existing grouping in a report.
	Matrix. Add or modify existing matrixes in a report.
	Copy to Group. Copy selected servers to a group.
	Font Size. Change font size (small, medium, or large).


Commands	Tooltip/Description
	Print/Quick Print. Show the Print Options dialog box to allow you to pick a printer and its settings, or just print the report immediately with the current default printer and settings.
	Search. Opens a Find dialog box for searching.
	Zoom Out, Zoom Percent, Zoom In. Zoom in or out incrementally/decrementally, or specify a zoom setting.
	Multiple Pages. Reduce the view zoom setting so you can view multiple pages of the report simultaneously.
	Color Background. Show a dialog box to allow you to set the background color of the report.
	Watermark. Show a dialog box where you can set a watermark for the report, such as DRAFT, TOP SECRET or an image.
	Export Document. Export to PDF, HTML, Text, CSV, MHT, Excel, RTF, or Image formats.
	Send via Email. Send the report via e-mail as an attachment. The file is saved locally first and then your default mail client is launched with the saved file attached. Possible file formats are PDF, MHT, RTF, Excel, CSV, Text, or an image format.
	Exit. Close the report preview page.


5.6.2 Customizing Reports

After a report is generated, it can be further adjusted and customized. This includes revising which columns to include in the report, changing the font, and adding or modifying grouping or matrixes. These customizations are displayed in the Document Viewer and are included in the report when it is exported, but are not saved to the template.

- ♦ [“Customizing Columns” on page 101](#)
- ♦ [“Changing Report Font Size” on page 102](#)
- ♦ [“Adding or Modifying Groups” on page 103](#)
- ♦ [“Adding or Modifying Matrixes” on page 103](#)

Customizing Columns

- 1 Click *Columns*  in the Document Viewer.
The Report Column Selection dialog box is displayed.



- 2 Select a column and use the arrow buttons to rearrange how it appears in the report (the higher the column in the tree hierarchy, the more to the left it appears in the report). Conditional formatting can be selected for numeric fields. Select custom colors and breakpoints, or use the default settings.

The report will highlight the following:


- ♦ All the values less than the *Low Breakpoint* with the *Low Color*
- ♦ All the values between the *Low* and *High Breakpoints* with the *Mid Color*
- ♦ All the values greater than the *High Breakpoint* with the *High Color*.

Selecting *Transparent* leaves the values without highlights.

- 3 Sort the report by selecting a criterion from the *Sort by* drop-down list. Select *Ascending* or *Descending* order. For grouped reports, this controls the order within the group, because the whole report is sorted by the grouping column.


Optionally, add summary rows to the table by selecting from the check boxes in the *Summary Rows* area. PlateSpin Recon calculates and displays the *Count*, *Sum*, *Maximum*, *Average*, or *Minimum* for numeric columns.

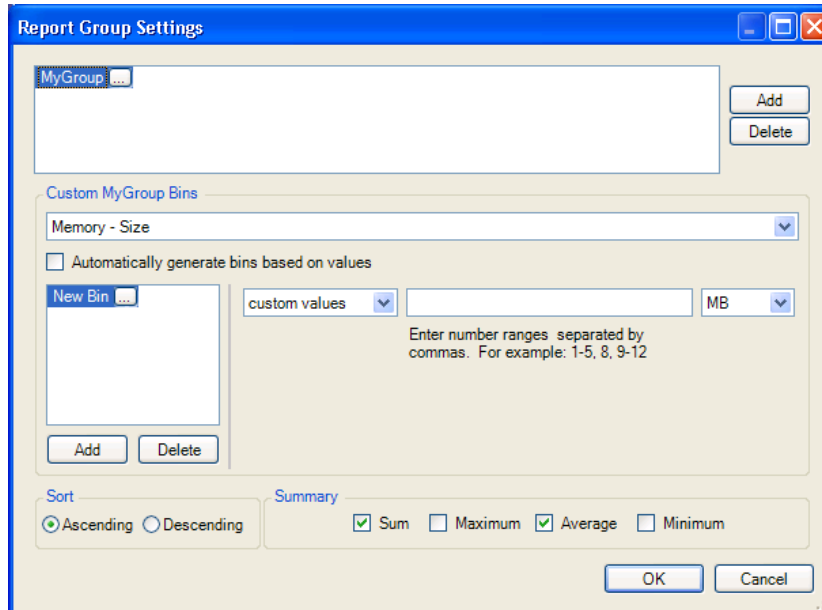
Changing Report Font Size

- 1 Click *Font*  in the Document Viewer. A drop-down menu is displayed, showing the following options :
 - ♦ Small Font
 - ♦ Medium Font
 - ♦ Large Font

Click the desired font size. The changes are applied to the report displayed in the Document Viewer.

Adding or Modifying Groups


- 1 Click *Groups*  in the Document Viewer.
The Report Group Settings dialog box is displayed.

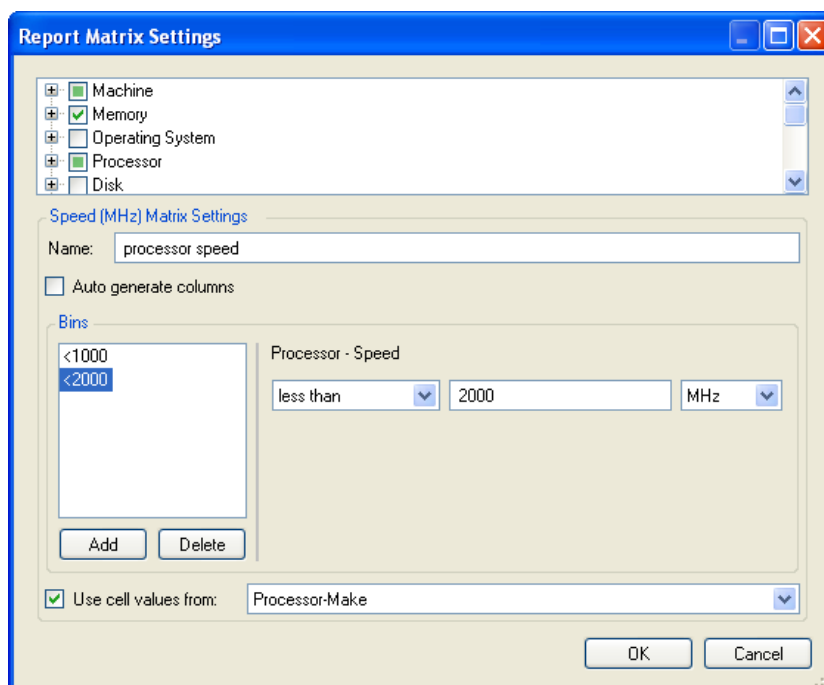


The **Report Group Settings** dialog box is shown. It has a title bar with standard window controls. Inside, there's a list box at the top containing 'MyGroup'. To its right are 'Add' and 'Delete' buttons. Below this is a section titled 'Custom MyGroup Bins'. It contains a dropdown menu set to 'Memory - Size'. Below the dropdown is an unchecked checkbox labeled 'Automatically generate bins based on values'. To the left of this checkbox is a 'New Bin' button and an empty list box. To the right of the checkbox is a text input field containing 'custom values', a unit dropdown set to 'MB', and a note: 'Enter number ranges separated by commas. For example: 1-5, 8, 9-12'. Below the 'New Bin' list box are 'Add' and 'Delete' buttons. At the bottom, there are two sections: 'Sort' with radio buttons for 'Ascending' (selected) and 'Descending', and 'Summary' with checkboxes for 'Sum' (checked), 'Maximum', 'Average' (checked), and 'Minimum'. 'OK' and 'Cancel' buttons are at the bottom right.

- 2 Click *Add* next to the upper list box to enter a group name.
- 3 In the Custom <groupname> Bins area, select the element to group by from the drop-down list.
- 4 Do either of the following:
 - ♦ Select the option *Automatically generate bins based on values*.
 - ♦ Click *Add* in the lower area to create and name a custom bin. Define an expression for each custom bin.
- 5 Select the desired sorting order for the group.
- 6 Select the summary rows to include in the table. These summarize the values for each bin.
- 7 To modify a group or a bin, select its name and edit it or its associated settings. Alternatively, click the appropriate *Delete* button to delete the group or bin.
- 8 Click *OK*.

Adding or Modifying Matrixes

- 1 Click *Matrix*  in the Document Viewer.
The Report Matrix Settings dialog box is displayed.



- 2 In the upper list box, select elements for creating a matrix.
- 3 For each selected element, enter a *Name* for its matrix.
- 4 Do either of the following:
 - ♦ Check *Auto generate columns* to create a column for each unique value of the element.
 - ♦ Click *Add* to add custom columns. For each column, specify a name and define an expression.
- 5 Check *Use cell values from* to specify a value to place in the appropriate matrix cells. If this box is left unselected, the values in the matrix columns are either True or False.
- 6 To remove a matrix from a report, deselect its check box. To modify a matrix, select it and modify its associated settings. Delete a bin by selecting its name and clicking *Delete*.


5.6.3 Viewing Reports

- ♦ “Diagnostics Reports” on page 104
- ♦ “Multiple Report Pages” on page 104



Diagnostics Reports

Click a row in a Missing Data - Summary report to view the corresponding Missing Data - Detailed and Error reports.

Multiple Report Pages

To view multiple pages at once, click *Multiple Pages* . Select the number of pages across and the number of pages vertically to view simultaneously. The more pages viewed, the smaller the text. The options available are from 1X1 through 5X10 pages.

5.6.4 Printing and Exporting Reports

Click *Print*  to open the Print dialog box, or click *Quick Print*  to immediately print the report, using the current default printer settings.



Click *Export Document*  to export data to any one of the formats, such as PDF, HTML, Text, CSV, MHT, Excel, RTF, or an image format. Select the desired format from the drop-down list. Click *Export Document* to automatically export data in the format most recently selected in the session. If a format has not yet been selected in this session, the default format is PDF.

Image Formats

After selecting to export a document as an *Image*, the Image Export Options dialog box is displayed.

Figure 5-11 Image Exports Options Dialog Box



Use this dialog box to select from several image types to export, including `bmp`, `emf`, `wmf`, `gif`, `jpg`, `png`, and `tiff`. You can also specify the following options to control the actions output:

Resolution (dpi): Image resolution you want to use, keeping in mind that the higher the resolution, the bigger the exported file.

Export mode: Allows you to choose how you want to export the image:

- ♦ **single file:** One long image with a single set of column headers at the top.
- ♦ **single file page-by-page:** One long image made up of individual pages, each with a set of column headers at the top.
- ♦ **different files:** A set of separate files, each with a single page image in it. When this option is selected, numbers are added to the end of the filename you specify for each separate file, incrementing until all the pages are exported. For example, if you specified `MyImage.bmp` for a two-page report, the two files exported would be named `MyImage01.bmp` and `MyImage02.bmp`.

Page range: Range of pages in the report to export. For example, to export the first five pages and the tenth page of a 20 page report, specify `1-5,10`.


Page border color: The color to use for the report page border.

Page border width: The width of the page border.

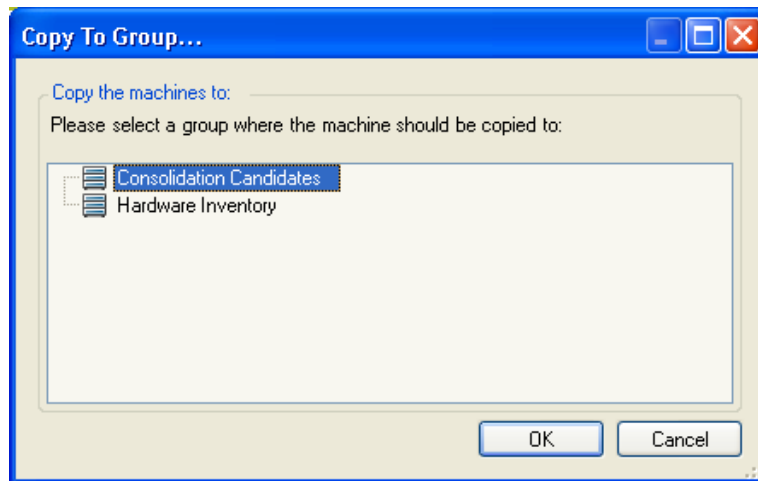
5.6.5 Copying Servers from a Report to a Group

Sometimes it is useful to copy the displayed servers from a report to a Data Center Explorer group. Not all reports display servers, but those types of reports do not display the Copy To Group tool. The groups must already exist. See [Section 3.4, “Working with Groups,” on page 29](#).

To copy servers to a group:

- 1 Click the check boxes next to the servers in the report you want to add to a group.
- 2 Click *Copy To Group* .

The Copy To Group dialog box is displayed.



- 3 Select the target group and click *OK*. To exit without copying, click *Cancel*.

Server Consolidation

6

This section provides information about the PlateSpin® Recon Server Consolidation and Protection features.

- ♦ [Section 6.1, “About Server Consolidation,” on page 107](#)
- ♦ [Section 6.2, “Project Explorer,” on page 107](#)
- ♦ [Section 6.3, “Consolidation Scenarios,” on page 120](#)
- ♦ [Section 6.4, “Generating Inventory Reports,” on page 131](#)
- ♦ [Section 6.5, “Viewing Consolidation Reports,” on page 132](#)
- ♦ [Section 6.6, “Implementing Scenarios with Portability Suite,” on page 136](#)

6.1 About Server Consolidation

The PlateSpin Recon Consolidation Planning Module allows for exploration of scale-up and scale-out consolidation scenarios with candidate workloads. Scenarios can assist with determining the minimum number of target servers and balancing workloads across these servers.

The key benefits of the Consolidation Planning Module include:

- ♦ Maximizing server utilization by distributing the number of workloads.
- ♦ Minimizing resource contention by balancing software workloads across five dimensions: Processor, Memory, Disk, Network, and Time.
- ♦ Eliminating the need for extra servers with advanced workload sizing and optimization.
- ♦ Identifying servers to be protected in a disaster recovery plan.

The following are the key stages to assessing server consolidation:

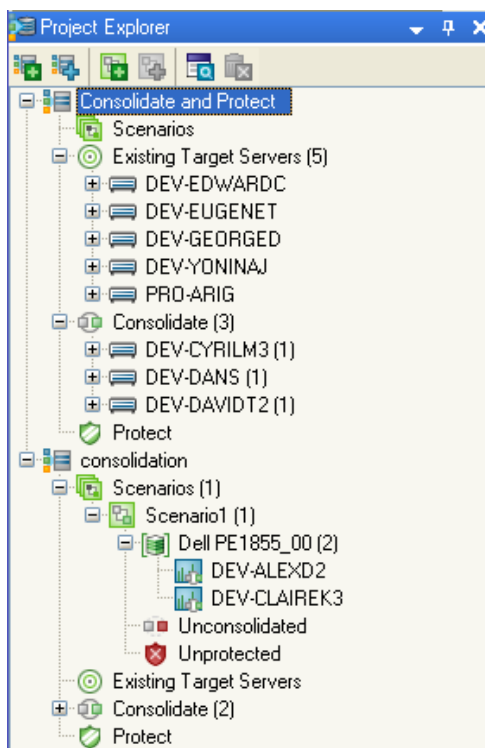
- 1 Discover and list servers in the PlateSpin Recon Client. For more information, see [Section 3.6, “Discovering and Inventorying Data Center Resources,” on page 32](#).
- 2 Inventory servers and collect static data. For more information, see [Section 3.6, “Discovering and Inventorying Data Center Resources,” on page 32](#).
- 3 Monitor servers and collect dynamic data. For more information, see [Section 3.11, “Viewing Inventory Data and Monitoring Servers,” on page 43](#).
- 4 Identify the optimal consolidation scenario. For more information, see [Section 6.3, “Consolidation Scenarios,” on page 120](#).
- 5 Select servers for consolidation. For more information, see [Section 6.3.5, “Identifying Consolidation Candidates,” on page 131](#).

6.2 Project Explorer

- ♦ [Section 6.2.1, “Creating Projects,” on page 109](#)
- ♦ [Section 6.2.2, “Server Templates,” on page 114](#)










Projects are displayed in the Project Explorer.

Figure 6-1 Typical Project Explorer Tree
















The project, scenario, and consolidated workloads are arranged hierarchically; that is, workloads are listed under a server, servers are listed under a scenario, and scenarios are listed under a project. All selected candidates are listed under workloads inside the project.

Table 6-1 Project Explorer Commands

Command	Action
	Create a new project.
	Copy a project
	View project settings
	Create a new scenario.
	Copy a scenario
	View report/View Project report/View Workload Assignment report.
	Delete project/scenario/workload, depending on the selection made.
	Implement by using the Portability Suite
	Edit Mapping

Alternatively, select Project Explorer commands by right-clicking any server, scenario, or project.

Table 6-2 *Project Explorer Icons*

Icon	Description
	Consolidation Project
	Consolidation Scenarios
	Server
	Existing Target Servers
	Virtual Machine
	Workloads
	Virtualization Overhead
	Server Template is mapped to a Virtual Server in the database
	Server is selected for Consolidation
	Unconsolidated Server
	Server is selected for Protection
	Server is Unprotected
	Unconsolidated Workloads

6.2.1 Creating Projects



A project can consist of any combination of consolidation candidates, protection candidates, workloads, and scenarios. Consolidation candidates are workloads on the network that you might want to move to virtual machines. Protection candidates are workloads you might want to create standby VM duplicates of for disaster recovery.

You create and compare scenarios in a project. The Project Report consists of scenario comparison results, which help in determining the best consolidation scenario for the environment.

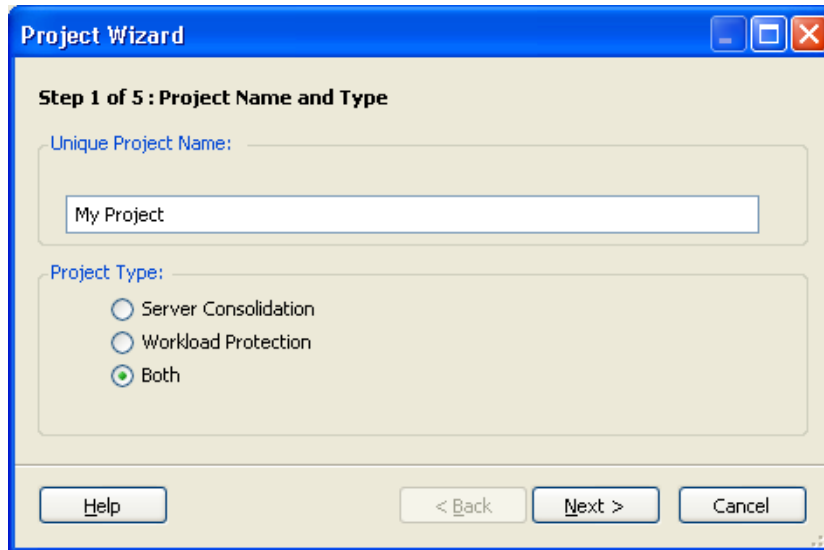
Before you create a project, identify and group consolidation candidates. The identified candidates allow you to create and compare consolidation scenarios in a project.

NOTE: Collecting at least 24 hours of data for each workload before starting the project is recommended.

To create a new project:

- 1 Click *Consolidation*  and select *New Project* or click *New Project*  in the Project Explorer.

The Project Wizard (Step 1 of 5: Project Name and Type) is displayed.



2 Type a *Unique Project Name* for this project in the text box.

3 Select a *Project Type*:

Server Consolidation: The project produces a recommendation for the best placement of workloads with respect to concurrent execution and CPU and memory headspace requirements.

Workload Protection: The project produces a recommendation for deployment of workloads for disaster recovery. Specify existing virtual servers or use server templates to build a scenario.

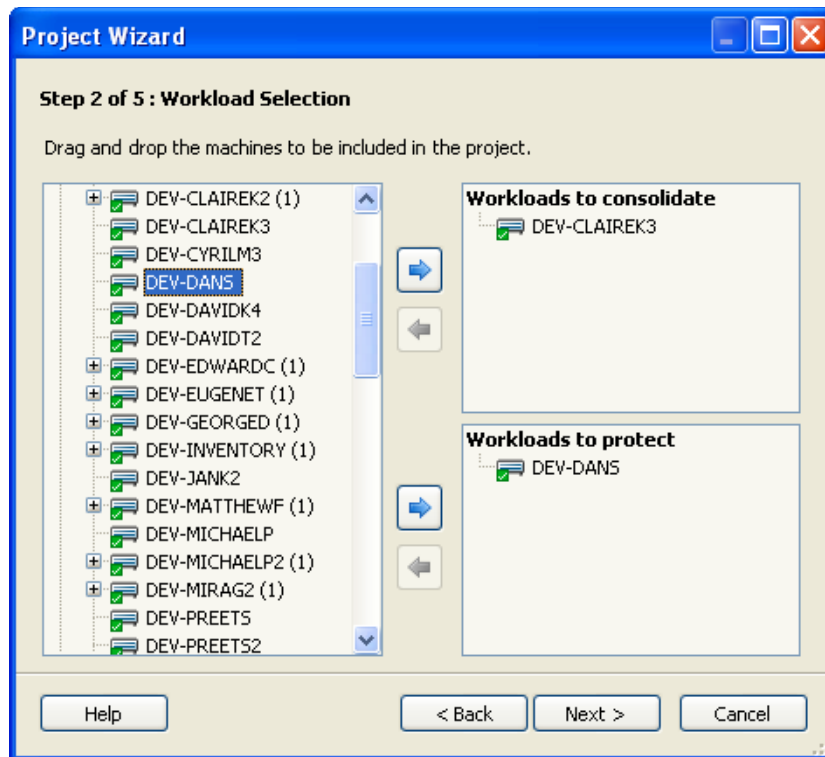
Both: The project produces a recommendation for both consolidation and protection. For example, optimally consolidate servers, and provision the remaining headroom for data protection.

A workload cannot be provisioned for both consolidation and protection in the same project.

4 Click *Next*.

The Step 2 of 5: Workload Selection dialog box is displayed.

5 Identify the workloads to include in the project.

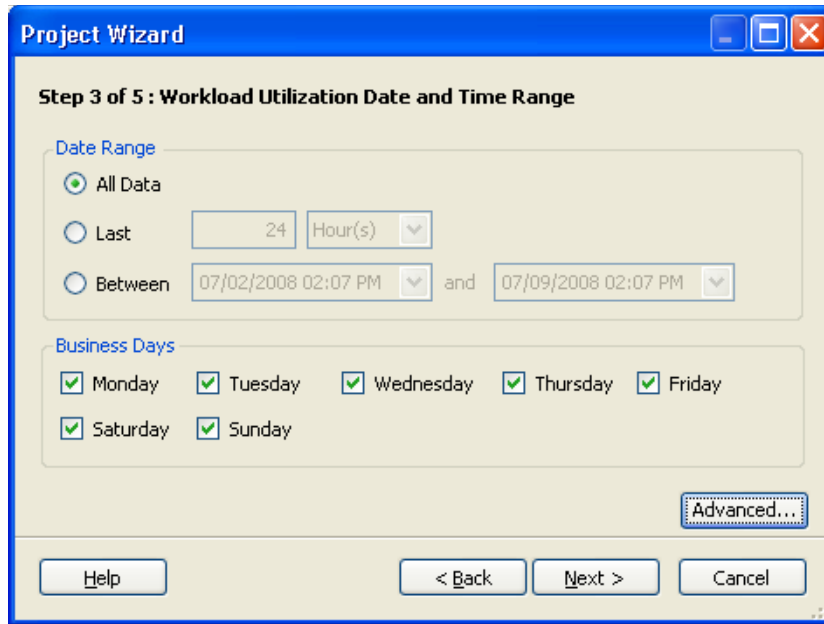


- 6 Drag workloads from the tree on the left and drop them into the Workloads to Consolidate panel or the Workloads to Protect panel on the right.

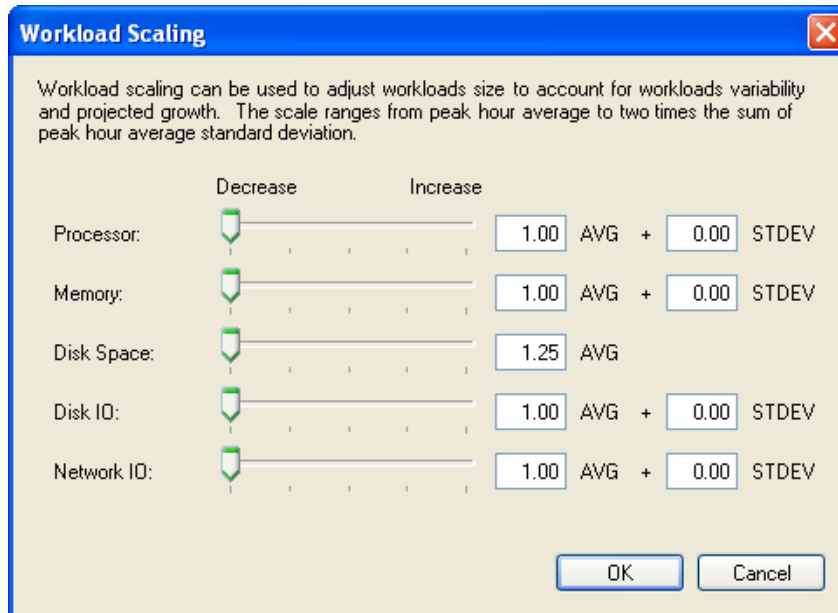
Alternatively, select servers and click the arrow pointing to the panel to which you want to add them. If you make a mistake, select the server in the panel on the right that you want to remove and click the arrow pointing back toward the server list. Servers cannot be targeted for consolidation and protection within the same project.

- 7 Drag and drop Target Servers into either of the two panels on the right.
- 8 Click *Next*.

The Step 3 of 5: Workload Utilization Date and Time Range dialog box is displayed.



- 9 Set the date and time range to summarize the workload data. *Date Range* and *Business Days* are required for calculating the 24H profile of workloads.
- 10 Optionally, click *Advanced* to manually set workload scaling factors.
The Workload Scaling dialog box is displayed.
- 11 Use the slide controls to adjust values or type the desired number.



Workloads can be scaled across five dimensions to account for workload variability and projected growth. Server resources allocated to workloads on target servers are based on scaled workloads. By scaling workloads upward, more server resources might be required, potentially resulting in lower consolidation ratios. When scaling pointers are moved, workload scaling factors change.

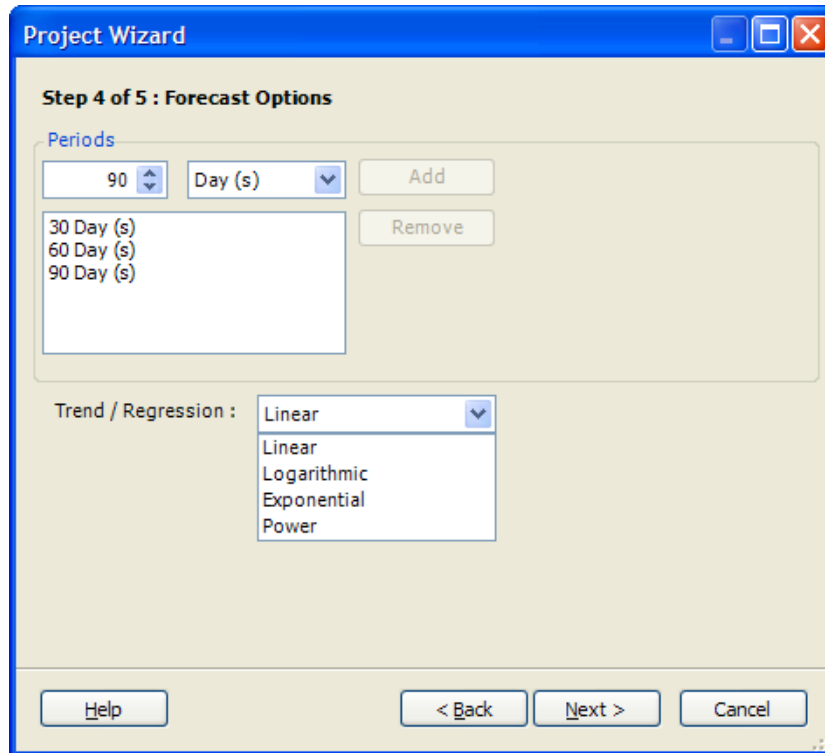
Workloads cannot be scaled below their peak hourly averages.

12 Click *OK*.

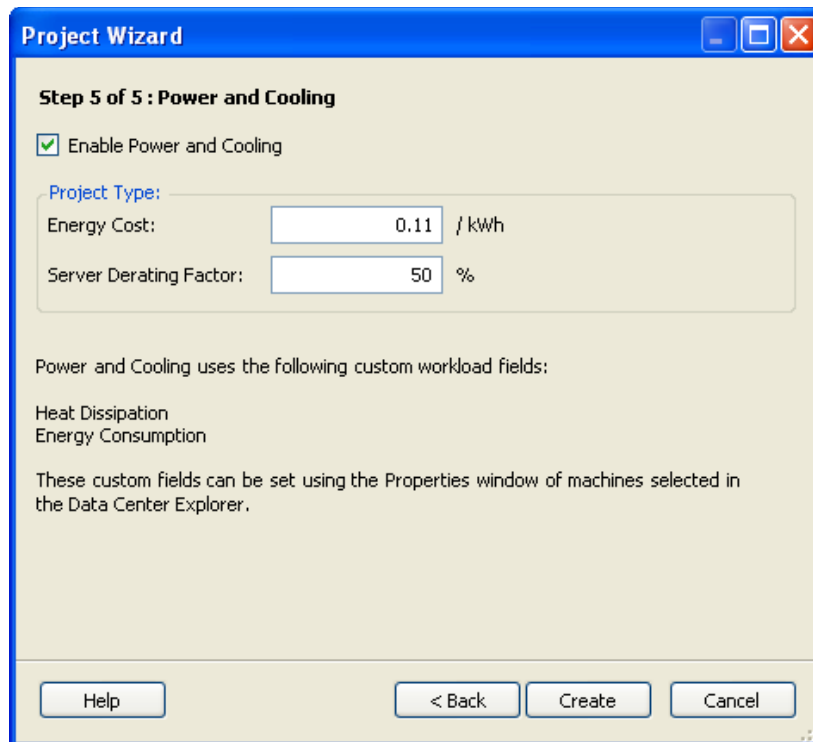
13 Click *Next*.

The Step 4 of 5: Forecast Options dialog box is displayed.

Select forecast options and a trend type. Workloads can be forecasted to account for future growth.



The Step 5 of 5: Power and Cooling dialog box is displayed.



14 Click *Next*.

15 Enter values for the *Project Goal*, or deselect the *Enable Power and Cooling* check box if these values are not to be included.

Energy Cost: Cost of energy per kilowatt hour.

Server Derating Factor: Percentage of total energy consumption by the average server. For example, a server rated at 800 watts uses 400 watts with a server derating factor of 50%.

16 Click *Create*.

The new project is displayed in the Project Explorer.




When creating a project, a 24H Profile of workloads is computed. This includes calculating Processor Speed (MHz), Memory (MB), Disk Space Used (GB), Disk IO (MB/s), and Network IO (MB/s) over a 24-hour period for each workload. The 24H Profile comprises 24 values of workloads that correspond to each hour of the day (0 to 23). Each value in the 24H Profile is calculated by taking the peak value from all values observed for that hour.

6.2.2 Server Templates

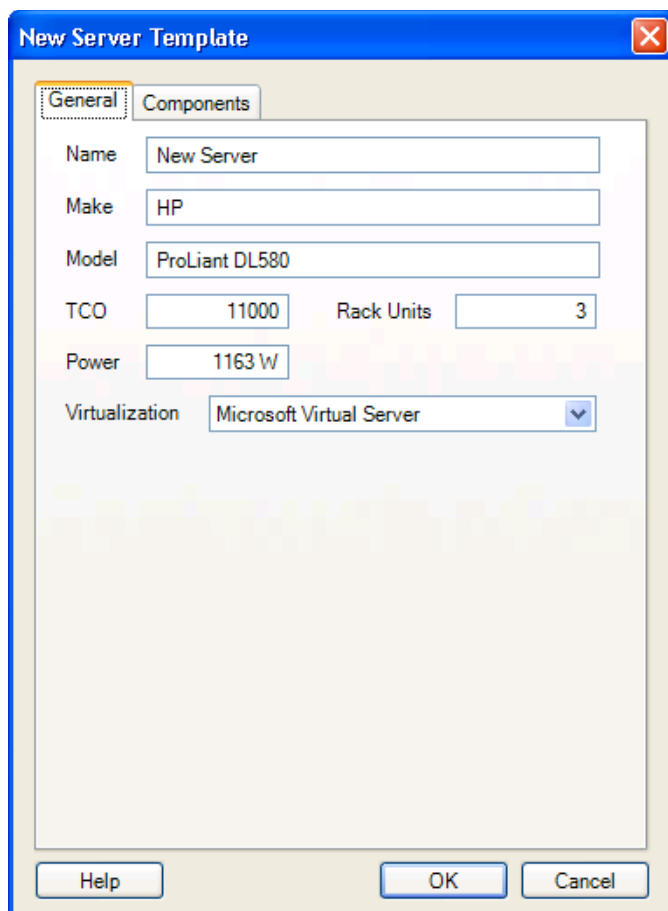
Server templates are used to define servers for use in consolidation scenarios. You create server templates, each with their own components and utilization thresholds. A common use for server templates is for scale-up or scale-out consolidation scenarios. The created server templates are used as target servers for consolidating workloads.

- ♦ “Creating a New Server Template” on page 115
- ♦ “Editing an Existing Server Template” on page 118
- ♦ “Copying a Server Template” on page 120

Creating a New Server Template

- 1 Click *Consolidation*  *Consolidation* ▼, then click *Server Templates*  on the menu.
The Server Template Manager dialog box is displayed.
- 2 Click *New* .

The New Server Template dialog box is displayed.



The image shows the 'New Server Template' dialog box with the 'General' tab selected. The fields are filled with the following values:

Field	Value
Name	New Server
Make	HP
Model	ProLiant DL580
TCO	11000
Rack Units	3
Power	1163 W
Virtualization	Microsoft Virtual Server

Buttons at the bottom: Help, OK, Cancel.

- 3 Make sure the General page is displayed, then specify the following :

User Interface	Required	Description
Name	Yes	Name of the template to be created.
Make	No	Server manufacturer.
Model	No	Server model.
TCO	No	Total cost of ownership.
Rack Units	No	Number of rack units.
Power	No	Power consumption in watts.

User Interface	Required	Description
Virtualization	Yes	Virtualization platform: <ul style="list-style-type: none"> ♦ Microsoft* Virtual Server ♦ Virtual Iron* ♦ VMware* ESX ♦ VMware ESX 3 ♦ VMware GSX ♦ VMware Server ♦ Citrix* XenServer* ♦ Microsoft Hyper-V* ♦ SLES Xen Server ♦ None

4 Click *Components*, then specify the following on the Components page:

New Server Template

General Components

Settings

Processor Count: 1

Processor Speed: 3333 MHz

Cores Per Processor: 1

CPU Model: Core Xeon

☐ Hyper-Threading

Normalized Speed: 41.10

Memory Size: 32768 MB

Disk Space: ☐ Fixed 0 GB ☒ Unlimited

Thresholds

Processor: 80 % Disk IO: 50 MB/s

Memory: 90 % Network IO: 10 MB/s

Disk Space: 85 %

Help OK Cancel

User Interface	Required	Description
Settings		

User Interface	Required	Description
Processor Count	Yes	Total processor count.
Processor Speed	Yes	Instructions per second the server can execute in MHz or GHz. The minimum processor speed you can use is 266 MHz.
Memory	Yes	Memory size for server in MB or GB.
Cores Per Processor	Yes	Number of cores per processor.
CPU Model	Yes	Model of CPU: <ul style="list-style-type: none"> ♦ Pentium* IV ♦ Core ♦ Xeon* ♦ Core Xeon ♦ Athlon* ♦ Opteron*
Hyper-Threading	No	Specifies if hyper-threading is available.
Normalized Speed	Readonly	A unitless measure of processor capacity. This provides more realistic comparisons across diverse processor architectures.
Memory Size	Yes	The amount of memory to use. The minimum memory size you can use is 256 MB.
Disk Space	Yes	Disk space in GB or TB that is required for the server. Select <i>Unlimited</i> if you are connecting target servers to a SAN or NAS, where disk space is not a limiting factor. If you use <i>Fixed</i> , the minimum amount of disk space you can specify is 1 GB.
Thresholds		
Processor	No	Processor utilization. By default, the utilization is set to 80%.
Memory	No	Memory utilization. By default, the utilization is set to 90%.
Disk Space	No	Maximum disk utilization threshold. By default, the utilization is set to 85%.
Disk IO	No	Maximum disk I/O allowed. By default, the Max IO is set to 50 MB/s.
Network IO	No	Maximum network IO. By default, the network IO is set to 10 MB/s.



Thresholds specify the level of utilization the Consolidation Planning Module can assume for the server templates. This affects the outcome of consolidation scenarios; lower utilization thresholds result in more servers being required to host the workloads of a scenario.

For example: workloads requiring 180 GB of memory can be consolidated onto a server with 200 GB of memory if the utilization threshold is set at 90%. However, if the memory utilization threshold is set at 50%, only 100 GB of memory can be used, resulting in fewer workloads being consolidated.

Utilization thresholds are initially set to the defaults listed in the table above and can be customized for specific needs.


- 5 Click *OK*.

Editing an Existing Server Template

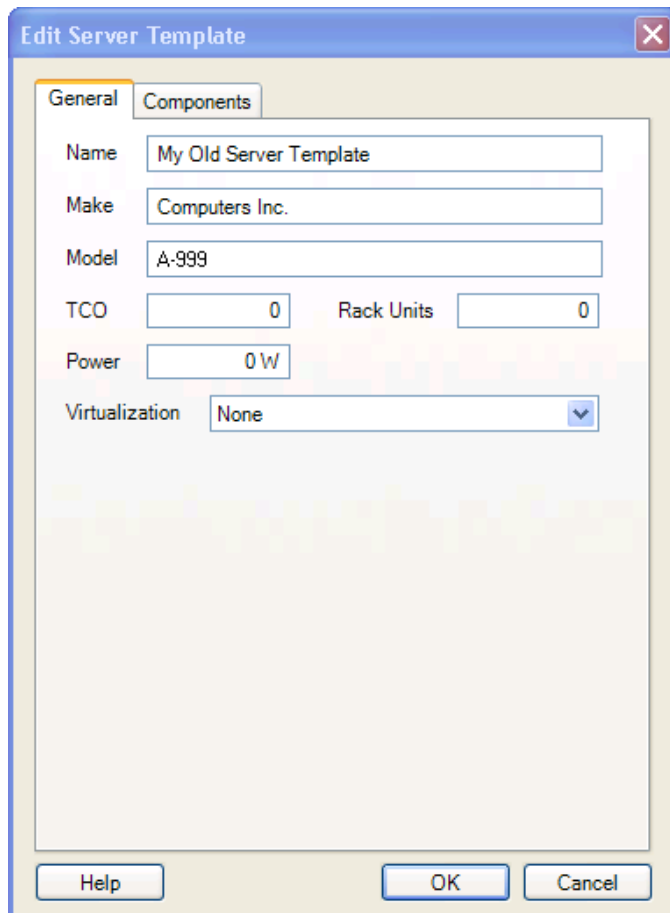
- 1 Click *Consolidation*  *Consolidation* ▼, then click *Server Templates*  on the menu.

The Server Template Manager dialog box is displayed.

- 2 Select the server template you want to edit.

- 3 Click *Edit* .

The Edit Server Template dialog box is displayed.



- 4 On the General page, make any desired modifications using the following as a guide :

User Interface	Required	Description
Name	Yes	Name of the template to be created.

User Interface	Required	Description
Make	No	Server manufacturer.
Model	No	Server model.
TCO	No	Total cost of ownership.
Rack Units	No	Number of rack units.
Power	No	Power consumption in watts.
Virtualization	Yes	Virtualization platform: <ul style="list-style-type: none"> ♦ Microsoft* Virtual Server ♦ Virtual Iron* ♦ VMware* ESX ♦ VMware ESX 3 ♦ VMware GSX ♦ VMware Server ♦ Citrix* XenServer* ♦ Microsoft Hyper-V* ♦ SLES Xen Server ♦ None

5 Click *Components*, then make any desired modifications using the following as a guide:

User Interface	Required	Description
Settings		
Processor Count	Yes	Total processor count.
Processor Speed	Yes	Instructions per second the server can execute in MHz or GHz. The minimum processor speed you can use is 266 MHz.
Memory	Yes	Memory size for server in MB or GB.
Cores Per Processor	Yes	Number of cores per processor.
CPU Model	Yes	Model of CPU: <ul style="list-style-type: none"> ♦ Pentium* IV ♦ Core ♦ Xeon* ♦ Core Xeon ♦ Athlon* ♦ Opteron*
Hyper-Threading	No	Specifies if hyper-threading is available.
Normalized Speed	Readonly	A unitless measure of processor capacity. This provides more realistic comparisons across diverse processor architectures.

User Interface	Required	Description
Memory Size	Yes	The amount of memory to use. The minimum memory size you can use is 256 MB.
Disk Space	Yes	Disk space in GB or TB that is required for the server. Select <i>Unlimited</i> if you are connecting target servers to a SAN or NAS, where disk space is not a limiting factor. If you use <i>Fixed</i> , the minimum amount of disk space you can specify is 1 GB.
Thresholds		
Processor	No	Processor utilization. By default, the utilization is set to 80%.
Memory	No	Memory utilization. By default, the utilization is set to 90%.
Disk Space	No	Maximum disk utilization threshold. By default, the utilization is set to 85%.
Disk IO	No	Maximum disk I/O allowed. By default, the Max IO is set to 50 MB/s.
Network IO	No	Maximum network IO. By default, the network IO is set to 10 MB/s.

Thresholds specify the level of utilization the Consolidation Planning Module can assume for the server templates. This affects the outcome of consolidation scenarios; lower utilization thresholds result in more servers being required to host the workloads of a scenario.

For example: workloads requiring 180 GB of memory can be consolidated onto a server with 200 GB of memory if the utilization threshold is set at 90%. However, if the memory utilization threshold is set at 50%, only 100 GB of memory can be used, resulting in fewer workloads being consolidated.

Utilization thresholds are initially set to the defaults listed in the table above and can be customized for specific needs.

6 Click *OK*.

Copying a Server Template

1 Select the template and click *Copy* .

The New Server Template dialog box is displayed with the template information prepopulated.

2 Rename the template and make any other required modifications.

3 Click *OK*.

The template is copied and is displayed in the Server Template Manager listing.

6.3 Consolidation Scenarios

Scenarios allow the quick creation and evaluation of one or more consolidation alternatives.

Generating multiple scenarios for comparison is recommended. This helps to determine the optimal consolidation scenario.


There are two types of scenarios:

- ♦ **Determine the minimum number of servers:** PlateSpin Recon determines how many of a specific type of target server are required to balance the workloads. With this type of scenario, only one server template can be used.
- ♦ **Specify the number of servers:** You set the number and type of target servers for consolidation. PlateSpin Recon calculates the optimal distribution of workloads for the chosen targets. With this type of scenario, multiple server templates can be used.

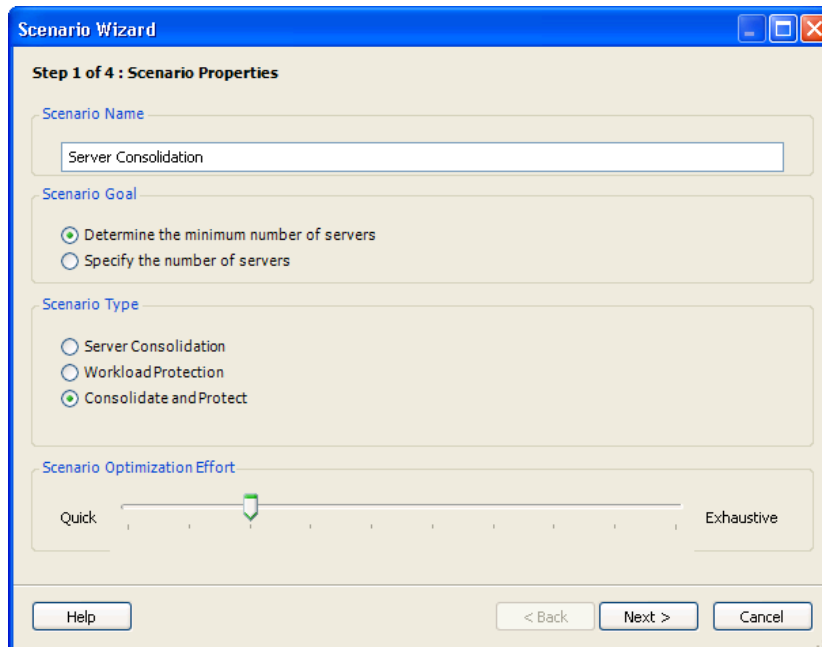
Any inventoried machine running any virtualization platform supported by PlateSpin Recon is considered a target server.

- ♦ [Section 6.3.1, “Determining the Minimum Number of Servers,” on page 121](#)
- ♦ [Section 6.3.2, “Specifying the Number of Servers,” on page 126](#)
- ♦ [Section 6.3.3, “Results of Consolidation Scenarios,” on page 129](#)
- ♦ [Section 6.3.4, “Editing Scenarios,” on page 130](#)
- ♦ [Section 6.3.5, “Identifying Consolidation Candidates,” on page 131](#)

6.3.1 Determining the Minimum Number of Servers

- 1 Select a project and click *New Scenario* .

The Scenario Wizard (Step 1 of 4: Scenario Properties) is displayed.



Available *Scenario Types* are determined by the *Project Goal* specified in the Project Wizard.

- 2 Specify a unique name in the *Scenario Name* text box.
- 3 Click *Determine the minimum number of servers*.

This goal determines the minimum number of a specific type of target servers to balance workloads. With this option, only one server template can be used.

4 Select a *Scenario Type*.

This step is relevant only if the project upon which the scenario is based specified *Both* (server consolidation and workload protection) as the *Project Goal* in the Project Wizard (see [Step 3 on page 121](#)).

5 Adjust the *Scenario Optimization Effort* slider. *Quick* is faster but not as intensive, and *Exhaustive* takes longer but is more thorough.

6 Click *Next*.

The Step 2 of 4: Workload Properties dialog box is displayed.

The screenshot shows the 'Scenario Wizard' dialog box, specifically 'Step 2 of 4: Workload Properties'. The dialog has a blue title bar and standard Windows window controls. It contains three main sections: 'Workload Sizing', 'Business Hours', and 'CPU Optimization Method'. In 'Workload Sizing', the 'Use historical size of workloads' radio button is selected, and the 'Specify forecast period' dropdown is set to '30 Day (s)'. In 'Business Hours', the 'Use the following hour range' checkbox is unchecked, and the time range is set from '09:00 AM' to '05:00 PM'. In 'CPU Optimization Method', the 'Use normalized speed to optimize workloads' radio button is selected. At the bottom, there are buttons for 'Help', '< Back', 'Next >', and 'Cancel'.

7 Specify Workload Sizing:

- ♦ **Use historical size of workloads:** Create the Scenario based on the historical size of the workloads.
- ♦ **Use forecasted size of workloads (specify forecast period):** Create the Scenario based on the forecasted size of the workloads after the specified forecast period.

For scenarios based on protection only projects, these values are dimmed .

8 To average workload and utilization expressions over a specified time range, select the *Use the following hour range* check box and set the start and end times for the business hours in your environment.

Hours specified here are used for all reports created for this scenario or for any node under this scenario. Report templates do not have any direct control over business/peak hours. Consolidation and Planning charts use both the specified business hours and all 24-hour data points.

9 Specify the CPU optimization method to use:

- ♦ **Use processor clock rate to optimize workloads:** Optimizes the CPU based on the actual processor clock.
- ♦ **Use normalized speed to optimize workloads:** Optimizes the CPU based on a normalized value. See [Section A.6, “CPU Normalization,” on page 162](#) for more information.

Consolidation charts based on this scenario will differ depending on your selection here. If you select normalized speed, the processor utilization and capacity values in consolidation charts created based on this scenario will also use the normalized values.

In any reports, Normalized Speed columns are not displayed by default, even if it's a scenario summary/workload assignment report and the optimization method selected is *Use normalized speed to optimize workloads*.

For scenarios based on protection only projects, these values are dimmed.

10 Click *Next*.

The Step 3 of 4: Workload Selection dialog box is displayed.

Scenario Wizard

Step 3 of 4 : Workload Selection


Select the workloads to be used in this scenario

Consolidate Protect

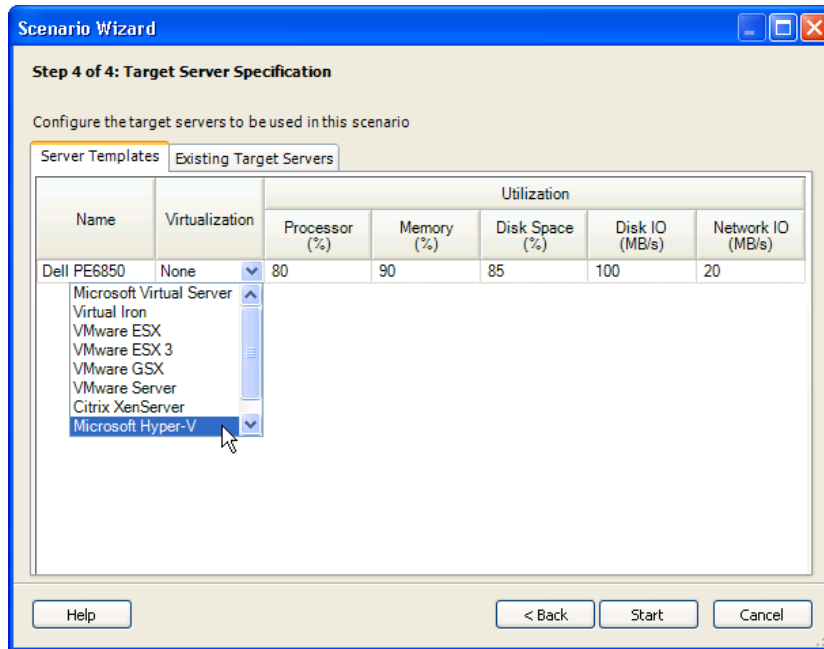
	Name	Processor (Normalized)	Memory (MB)	Disk Space (GB)	Disk IO (MB/s)	Network IO (MB/s)
<input checked="" type="checkbox"/>	DEV-HERBB	0.17	676.73	75.22	0.01	0.12
<input checked="" type="checkbox"/>	DEV-MARTINS	3.03	1430.31	36.37	0.96	0

Help < Back Next > Cancel

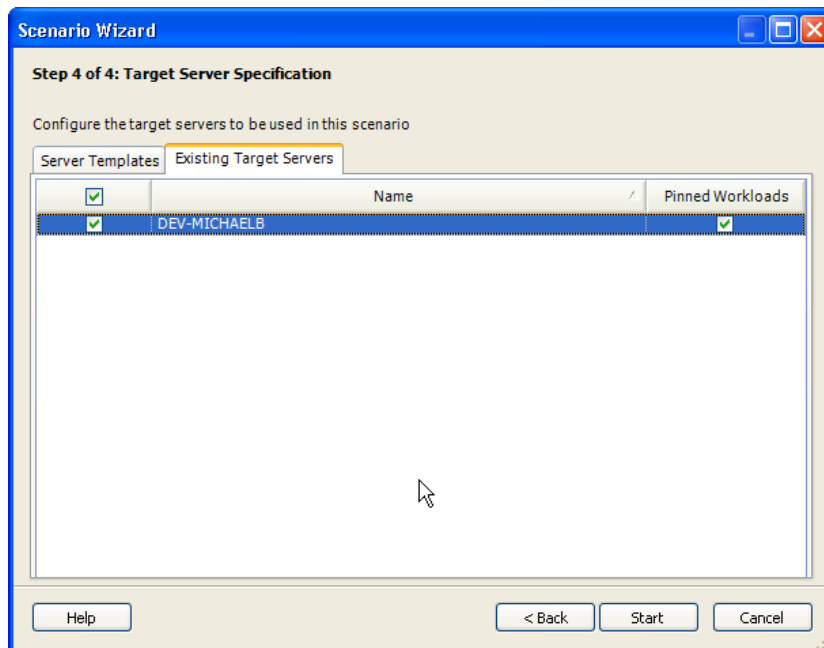
11 Select workloads or servers, depending on the scenario type.



- 12** (Conditional) If the *Scenario Type* includes server consolidation, select the servers to be consolidated in this scenario. By default, all servers specified in the project for consolidation are selected.
- 13** (Conditional) If the *Scenario Type* includes workload protection, select workloads on the *Protect* tab. By default, all workloads specified in the project for protection are selected.
- 14** Click *Next*.
The Step 4 of 4: Target Server Specification dialog box is displayed.




- 15 Select the *Server Template* to use from the *Name* drop-down list or create a new server template by selecting *New Server*. In this scenario, one server template can be used.
You must specify at least one server template, even if your intent is to consolidate only to existing target servers.
- 16 Optional Click the *Virtualization* drop-down list and select the Virtualization platform to use.
- 17 (Optional) Click the *Existing Target Servers* tab to use servers included in the project. If there are no existing target servers specified in the project, this option is unavailable.

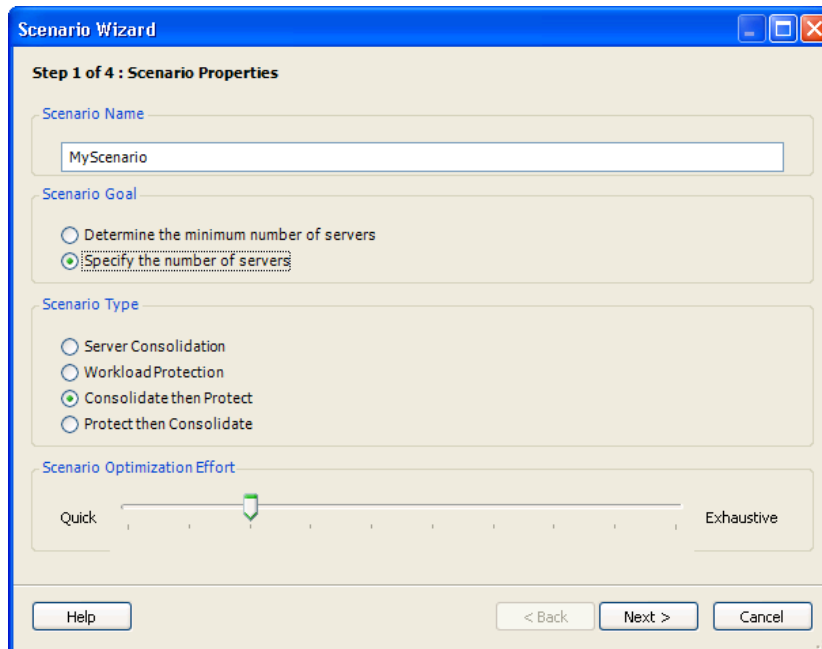


- 18 Select the existing target servers to use in the scenario, then select the *Pinned Workloads* check box to force workloads currently on a particular target server to remain on that server after consolidation.
- 19 Click *Start*.

The Create Scenario job runs. You can watch the progress in the Job Explorer window. When the scenario is created, it is displayed under Scenarios under the base project.

6.3.2 Specifying the Number of Servers


- 1 Click *New Scenario*  to launch the Scenario Wizard (Step 1 of 4: Scenario Properties).



- 2 Type a unique *Scenario Name*.
- 3 Select *Specify the number of servers* to set the number and type of target servers for consolidation. With this option, multiple server templates can be used.
- 4 Select the *Scenario Type*.

If some workloads must be left unconsolidated because the number of templates specified is inadequate, the algorithm fits as many of the consolidation workloads onto those templates as it can if you choose *Consolidate then Protect*. The reverse is true if you choose *Protect then Consolidate*.
- 5 Adjust the *Scenario Optimization Effort* slider. *Quick* is faster but not as intensive, and *Exhaustive* takes longer but is more thorough.
- 6 Click *Next*.

The Step 2 of 4: Workload Properties dialog box is displayed.



7 Specify Workload Sizing:

- ♦ **Use historical size of workloads:** Create the Scenario based on the historical size of the workloads.
- ♦ **Use forecasted size of workloads (specify forecast period):** Create the Scenario based on the forecasted size of the workloads after the specified forecast period.

For scenarios based on protection only projects, these values are dimmed.

8 To average workload and utilization expressions over a specified time range, select the *Use the following hour range* check box and set the start and end times for the business hours in your environment.

Hours specified here are used for all reports created for this scenario or for any node under this scenario. Report templates do not have any direct control over business/peak hours. Consolidation and Planning charts use both the specified business hours and all 24-hour data points.

9 Specify the CPU optimization method to use:

- 9a Use processor clock rate to optimize workloads:** Optimizes the CPU based on the actual processor clock. See [Section A.6, “CPU Normalization,” on page 162](#) for more information.
- 9b Use normalized speed to optimize workloads:** Optimizes the CPU based on a normalized value.


Consolidation charts based on this scenario will differ, depending on your selection here. If you select normalized speed, the processor utilization and capacity values in consolidation charts created based on this scenario will also use the normalized values.

In any reports, Normalized Speed related columns are not displayed by default, even if it's a scenario summary/workload assignment report and the optimization method selected is *Use normalized speed to optimize workloads*.

For scenarios based on protection only projects, these values are dimmed.

10 Click *Next*.

The Step 3 of 4: Workload Selection dialog box is displayed.



11 (Conditional) If the *Scenario Type* includes server consolidation, select the servers to be consolidated in this scenario. By default, all servers specified in the project for consolidation are selected.

12 (Conditional) If the *Scenario Type* includes workload protection, select workloads on the *Protect* tab. By default, all workloads specified in the project for protection are selected.

If you are running a scenario with a set number of target servers, in some cases the set number of target servers do not have sufficient resources to accommodate all workloads.

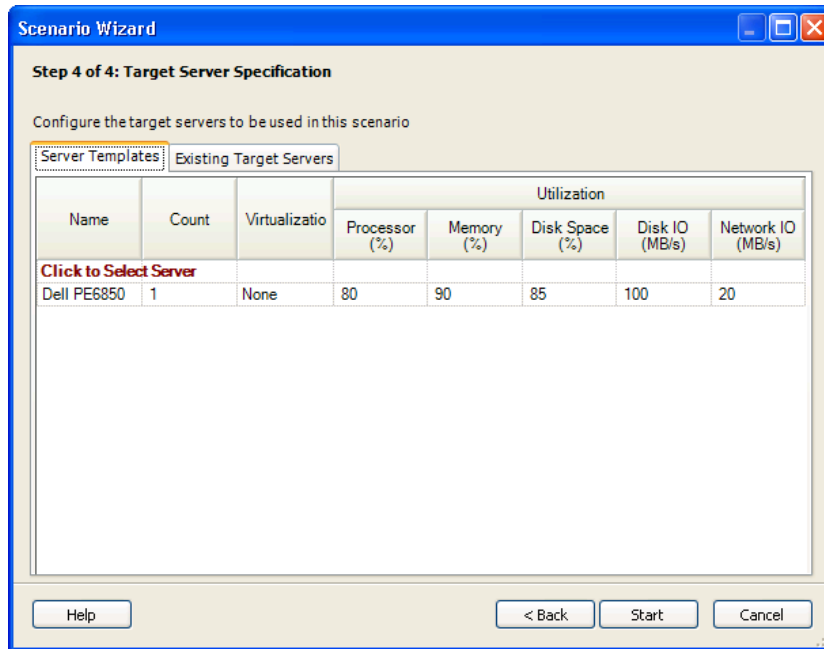
By default, *Honor utilization thresholds* is selected. This maintains the total utilization of a target server below the set thresholds. In other words, the sum of workloads will not violate the utilization thresholds for a specified server.

13 Define consolidation priorities by selecting one of these options:

- ♦ **Maximize the number of consolidated workloads:** Workloads are selected for consolidation with the goal of consolidating the largest possible number of workloads. The order shown is not relevant and the up and down arrows are disabled.
- ♦ **Consolidate workloads in the above order:** Prioritize workloads for consolidation only, in the given order. This ensures that certain workloads are consolidated. To move workloads into the desired order of consolidation priority, use the up and down arrows.

14 Click *Next*.

The Step 4 of 4: Target Server Specification dialog box is displayed.




- 15 Select the *Server Template* to use from the *Name* drop-down list or create a new server templates by selecting *Click to Select Server*.
- 16 Specify the *Count* for each server template.
- 17 It is possible to use existing target servers if they were included in the project. Click the *Existing Target Servers* tab to use servers specified in the project. If no existing target servers are specified in the project, this option is not available.
- 18 Click *Start*.

The Create Scenario job runs. You can watch the progress in the Job Explorer window. When the scenario is created, it is displayed under Scenarios under the base project.

6.3.3 Results of Consolidation Scenarios

The result of any scenario depends on the workloads used and the server templates selected. Higher consolidation ratios are achieved with larger target servers. The created scenario is displayed in the Project Explorer. For a description of the icons in the graphic below, see the listing of Project Explorer icons in [Table 6-2 on page 109](#).

Figure 6-2 Project Explorer




Virtualization Overhead: If server templates configured with a virtualization platform are selected as input to the scenario, the overhead related to the selected virtualization platform is shown as an additional workload under each target server and is named after the virtualization platform. Virtualization overhead workloads are not movable.

Unconsolidated Workloads: Lists all servers that were not consolidated during the scenario creation process. Workloads are rejected if they cannot be added to a server without exceeding utilization thresholds. These servers are listed under *Unconsolidated* in the scenario. Moving unconsolidated workloads to a server can result in a violation of utilization thresholds.

Unprotected Workloads: Lists all servers that were not protected during the scenario creation. Workloads are rejected if they cannot be added to a server without exceeding utilization thresholds. These servers are listed under *Unprotected* in the scenario. Moving unprotected workloads to a server can result in a violation of utilization thresholds.

6.3.4 Editing Scenarios

Move workloads between target servers or from *Unconsolidated* or *Unprotected* to a target server by dragging and dropping.

Select a scenario and click *Delete Scenario*  to remove consolidated workloads from target servers and move them to *Unconsolidated*. Alternatively, right-click the workload and click *Unconsolidate Workload* on the menu. Deleting protected workloads moves them to *Unprotected*.

6.3.5 Identifying Consolidation Candidates

This activity involves selecting servers for consolidation. This requires you to generate inventory and analysis reports so you can find a set of servers that meet your consolidation requirements, then you can select your consolidation candidates from that set of servers. For more information, see [Section 5.5, “Generating Reports,” on page 96](#).

6.4 Generating Inventory Reports

- ♦ [Section 6.4.1, “Finding Servers That Meet the Consolidation Criteria,” on page 131](#)
- ♦ [Section 6.4.2, “Selecting Candidates for Consolidation,” on page 131](#)

6.4.1 Finding Servers That Meet the Consolidation Criteria

- 1 In the Report Explorer, right-click the Hardware Inventory report template under *Inventory*, then select *Edit Report Template*.
The Edit Report Template dialog box is displayed.
- 2 On the *Group* tab, select a filter or group to use when you generate reports based on this template.
- 3 Click the *Views* tab and select a view. Use the buttons adjacent to the list to add or edit a view. See [Section 5.4, “View Customization,” on page 92](#)
- 4 Click the *Field Filter* tab. Use the drop-down lists to specify query parameters. See [“Field Filter Tab” on page 85](#).
- 5 Click *OK*.
- 6 In the Report Explorer, right-click *Hardware Inventory* and select *Generate Report*.
- 7 Schedule the report for *Now*, and accept the default *Delivery Type* of the *Report Explorer*.
- 8 Click *Generate*.

The report instance, shown as the date and time the report was generated, is displayed under the report template in the Report Explorer and the report is previewed in the Document Viewer, showing the servers that meet your consolidation requirements.

6.4.2 Selecting Candidates for Consolidation

- 1 In the Report Explorer, right-click the *Analysis* report template for Hardware Utilization - Peak and select *Edit Report Template* on the menu.
The Edit Report Template dialog box is displayed.
- 2 On the *Group* tab, make the same selection you made in [Step 1 on page 131](#).
- 3 On the *Views* tab, select the table view and click *Edit*.
The Create View dialog box is displayed.
- 4 On the *General* tab, select the *Summary Rows* (applied to columns containing calculated data) you want to create in the report.
- 5 Click *Finish*.

The Edit Report Template dialog box is displayed again.

- 6 On the *Date Filter* tab, set the date and time range to view utilization expressions for all servers that were monitored.
- 7 On the *Field Filter* tab, use the drop-down lists to specify the same query parameters you selected in [Step 1 on page 131](#).
- 8 On the *Summaries* tab, select *AVG*.
- 9 Click the *Calculation type* drop-down list and a type:
 - ♦ **Average:** Average of all hourly values across business days and business hours.
 - ♦ **Peak Hour:** A calculation based on the highest hourly value.
 - ♦ **Total:** Sums all the data points in the time frame. For example, Disk IOs per second for a 1-hour period equates to 10 disk IOs per second X 60 seconds X 60 minutes = 36000. Select the *Summary Rows* to create in the report of the columns containing calculated data.
- 10 Click *OK*.
- 11 In the Report Explorer, under Analysis, right-click the previously edited Hardware Utilization - Peak report template and click *Generate Report* on the menu. Accept the default *Delivery Type* of the *Report Explorer*.
- 12 Click *Generate*.

The report instance, shown as the date and time the report was generated, is displayed under the report template in the Report Explorer and the report is previewed in the Document Viewer.

Comparing and contrasting the two reports generated in [Step 1 on page 131](#) and [Step 1 on page 131](#) will help you determine your consolidation needs.

To help with this, while viewing a generated report, click a cell to display a customizable chart for that server/workload metric. For more information see [Section 5.1, “Chart Viewer,” on page 75](#).

6.5 Viewing Consolidation Reports

Consolidation reports are used for analysis. Unlike other PlateSpin Recon reports, which are generated through the Report Explorer, consolidation reports are generated through the Project Explorer. Like other PlateSpin Recon reports, generated consolidation reports are viewed in the Document Viewer. In addition, consolidation reports cannot be saved and must be generated whenever necessary.

The following types of consolidation reports are available:

- ♦ Project Report
- ♦ Scenario Summary Report
- ♦ Workload Assignment Report
- ♦ Unconsolidated Report
- ♦ Unprotected Report

Project, scenario summary, and workload assignment reports are template based. Edit the templates through the Report Explorer. For information about editing consolidation report templates, see [Section 5.3.3, “Consolidation Report Templates,” on page 89](#).

- ♦ [Section 6.5.1, “Project Report,” on page 133](#)
- ♦ [Section 6.5.2, “Scenario Summary Report,” on page 133](#)

- ♦ Section 6.5.3, “Workload Assignment Report,” on page 135
- ♦ Section 6.5.4, “Unconsolidated Workload Report,” on page 136
- ♦ Section 6.5.5, “Unprotected Workload Report,” on page 136

6.5.1 Project Report

A Project Report provides a comparison of consolidation scenarios. To view this report, double-click a project.

The Project Report is displayed in the Document Viewer.

Figure 6-3 Sample Project Report

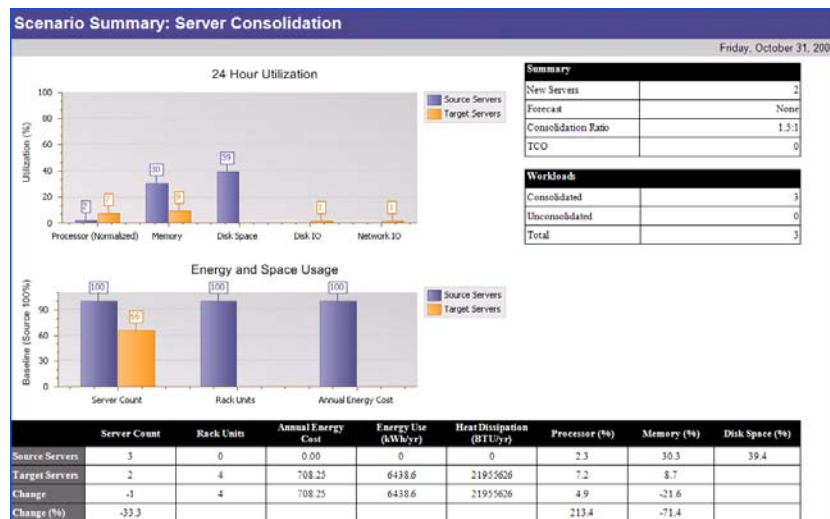
Project: Consolidation and Protection											
Friday, October 31, 2008											
Scenario	Utilization (%)			Workloads			New Servers				
Name	Processor (Spec)	Processor (Normalized)	Memory	Consolidated	Protected	Count	Consolidation Ratio	Protection Ratio	TCD	Rack Units	Power (W)
Before Consolidation	2.41	2.30	30.32	0	0	—	—:1	—	—	—	—
Server Consolidation	11.11	7.21	8.67	2	1	2	1.00:1:1	0.5:1	0	4	1470
Summary Statistics											
Max	11.11	7.21	8.67	2	1	2	1.00:1	0.5:1	0	4	1470
Avg	11.11	7.21	8.67	2	1	2	1.00:1	0.5:1	0	4	1470
Min	11.11	7.21	8.67	2	1	2	1.00:1	0.5:1	0	4	1470

6.5.2 Scenario Summary Report

The Scenario Summary Report helps in assessing the value of using a particular infrastructure for server consolidation.

Double-click the selected scenario to view the Scenario Summary report in the Document Viewer and the 24H Profile Stacked Utilization chart in the Chart Viewer.

Figure 6-4 Sample Scenario Summary Report



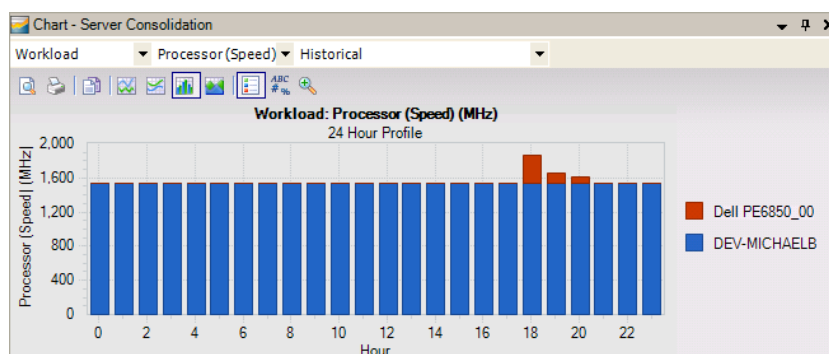
If there are multiple target servers, a summary statistics table is displayed with SUM, AVG, MAX, and MIN values.

Figure 6-5 Sample Report Summary Statistics Table

Machine		Processor	Memory	Disk		Network
Name	Workloads	% Mhz used (%)	Used (%)	Used (%)	IO (%)	IO (%)
DEV-MICHAELB	0	24.0	27.0	--	0.0	0.0
Dell PE6850_00	3	3.2	7.5	--	1.0	1.0
Summary Statistics						
Sum	3					
Max	3	24.0	27.0	--	1.0	1.0
Avg	2	13.6	17.3	--	0.5	0.5
Min	0	3.2	7.5	--	0.0	0.0

If you click one of the server cells, the corresponding 24H Profile stack chart is displayed, showing all workloads assigned to the target server.


Figure 6-6 Sample 24H Profile Stack Chart



24H Profile Chart Appearance

Twenty-four hour (24H) Profile charts in the Data Center Explorer and in Consolidation Planning are slightly different. In the Data Center Explorer, the 24H profile that is shown. select several monitored machines, click *Create Chart* and then select *Profile* as the type of chart) in workload profile charts is for each individual resource instance separately and shows a different color for each instance and resource. These values are stacked for each given hour. This gives the maximum value for every instance for a specific hour over the monitoring period. These stacked values are then summed.


Figure 6-7 Summed Stack Values



This highlights the hour with the maximum utilization for a specific instance of the resource. This approach is particularly useful for determining maximum utilization of any individual resource instance.

In Consolidation Planning, the 24H chart has the values of the instances and resources summed for each time moment and then the maximum of this sum for every hour is identified.

Figure 6-8 Consolidation Planning 24H Chart



This is needed to determine the summary resource consumption of a workload.

6.5.3 Workload Assignment Report

In the Workload Assignment report, you can see which workloads are assigned to servers and how workloads are utilized in each of these assigned servers. Then, you can modify workload assignments according to requirements and consider different kinds of server options while creating a project.

Workload Assignment reports are available for individual target servers or for all target servers included in a scenario.

- ♦ “Individual Server” on page 135
- ♦ “All Servers” on page 136

Individual Server

Double-click a target server in the *Project Explorer* to view the Workload Assignment report and 24H Profile chart.

NOTE: A 24H Profile cannot be created for workloads when there is no one-hour summary data.

If you click one of the servers in the Assignment report, the corresponding 24H Profile stack chart is displayed, showing all workloads assigned to the target server. The utilization values displayed are for the target servers.

Figure 6-9 Assignment Report Generated 24H Profile Stack Chart



All Servers

Right-click a scenario and select *View Report > View Workload Assignment Report* from the drop-down menu. The Workload Assignment report is displayed, showing data for all the servers in the current scenario.

6.5.4 Unconsolidated Workload Report

In the Unconsolidated Workload Report, you can see which of the consolidation candidate workloads are not included in the consolidation plan, as well as an explanation of the reason why PlateSpin Recon could not include them. To view this report, right-click a scenario and select *View Report > View Unconsolidated Report*. The Unconsolidated Report is displayed in the Document Viewer.

Figure 6-10 Sample Unconsolidated Report

Unconsolidated Report		
B384 Project > B384 Scenario2 > Unconsolidated		April 6, 2007
	Machine	
<input type="checkbox"/>	Machine Name	Reason
<input type="checkbox"/>	DEV-DAVIDT	There are not enough resources in any of the available containers
<input type="checkbox"/>	DEV-NADERS	There are not enough resources in any of the available containers
<input type="checkbox"/>	DEV-BORISB	There are not enough resources in any of the available containers
<input type="checkbox"/>	DEV-DAVIDK	There are not enough resources in any of the available containers

6.5.5 Unprotected Workload Report

In the Unprotected Report, you can see which of the consolidation candidate workloads were not included in the protection plan, as well as an explanation of the reason why PlateSpin Recon could not include them. To view this report, right-click a scenario and select *View Report > View Unprotected Report*. The Unconsolidated Report is displayed in the Document Viewer.

Figure 6-11 Sample Unprotected Report

Unprotected Report		
Project 2 Protect > Scenario A > Unprotected		Monday, April 09, 2007
	Machine	
<input type="checkbox"/>	Machine Name	Reason
<input type="checkbox"/>	DEV-DSBUILD	There are not enough resources in any of the available containers
<input type="checkbox"/>	FOGBUGZ	There are not enough resources in any of the available containers
<input type="checkbox"/>	VM-SRC3	There are not enough resources in any of the available containers
<input type="checkbox"/>	VM-SRC6	There are not enough resources in any of the available containers
<input type="checkbox"/>	YONINAVM	There are not enough resources in any of the available containers

6.6 Implementing Scenarios with Portability Suite

When there are identified scenarios and servers to deploy, you can implement the consolidation and protection by using PlateSpin Portability Suite. You can implement an entire scenario, separate containers within the scenario, or individual servers in the scenario.

To implement a scenario:

- 1 Right-click the scenario and select *Implement using Portability Suite*.

The Implementing scenario <name> using Portability Suite Step 1 of 2 dialog box is displayed.

Implementing scenario MyScenario (2) using PortabilitySuite

Step 1 of 2:

PortabilitySuite Settings

Specify the PortabilitySuite Server you will be using for the consolidation.

PortabilitySuite Server URL:

Network Name:

PortabilitySuite Credential:

Destination

Specify the folder relative to the server where the PortabilitySuite jobs should be saved. Note that the folder should be accessible by the account assigned to the PlateSpin Recon Service.

Directory:

Help < Back Next > Cancel

- 2 Confirm the *Portability Suite Server URL* and *Network name*.
- 3 Use the drop-down list to select or create the credentials that Portability Suite uses to discover and convert the servers within the scenario or container.
- 4 Specify the folder where the Portability Suite jobs should be saved.
- 5 Click *Next*.

The Step 2 of 2 dialog box is displayed.

Implementing scenario MyScenario (2) using PortabilitySuite

Step 2 of 2:

Map the containers to physical servers by selecting them in the list below.

Container	Server
Dell PE6850_00	Not Mapped
DEV-MICHAELB	Not Mapped

Help < Back Start Cancel

- 6 Click *Not Mapped*. A table drops down where a match column is displayed.

If you are attempting to consolidate multiple containers, the mapping of container to target server is one-to-one. Unavailable containers that are already mapped to a server are dimmed . When they are either unmapped or the scenario they are in is deleted, they are once again made available.

7 Click *Start*.

When the job is initiated, PlateSpin Recon:

- ♦ Authenticates the servers.
- ♦ Creates and runs server discoveries for Portability Suite.
- ♦ Creates and saves conversion jobs in the file location specified in [Step 1 on page 137](#). The default location is `c:\temp`.

When the job completes, you can load and execute the conversion jobs in Portability Suite.

6.6.1 Implementing Phantom or Existing Target Servers

The steps to implement at the “phantom” (server template), or existing target server level, are identical to implementing a scenario: right-click the phantom or existing target virtual machine server and select *Implement Using Portability Suite*.

If a phantom server represents an existing target virtual machine server, you can map it to that existing server. When the job is created through Portability Suite, the source and target servers are auto-configured.

To map a phantom to a target virtual machine server and deploy it:

1 Right-click the phantom and select *Edit Mapping*.

The Server Mapping dialog box is displayed, listing the available servers that contain virtual machines. Data is listed for each machine, and the match column indicates whether or not the machine is a match to the phantom.

Test	Results
CPU	<p>Failure: The processing power of the destination server is less than the processing power of the source server.</p> <p>Warning: The destination server has more processing power than is necessary. A warning is displayed if the processing power of the destination server is 500 Mhz greater than the processing power of the source server.</p> <p>Pass: The CPU match is acceptable.</p>
Disk	<p>Failure: The disk space on the destination server is less than the disk space on the source server.</p> <p>Warning: The disk space on the destination server is higher than necessary. A warning is displayed if the destination server has more than 50 GB of surplus space compared to the source server.</p> <p>Pass: The disk match is acceptable.</p>

Test	Results
Memory	<p>Failure: The destination server has less memory than the source server.</p> <p>Warning: The destination server has more memory than is necessary. A warning is displayed if the memory on the destination server is 512 MB greater than the memory on the source server.</p> <p>Pass: The memory match is acceptable.</p>

- 2 In the Server Mapping dialog box, select the target virtual machine server to map the phantom to, then click *OK*.
- 3 Right-click the newly mapped phantom and select *Implement Using Portability Suite*.
- 4 Complete the Implementing Server Using Portability Suite Wizard as described in “**To implement a scenario:**” on page 137.
- 5 Click *Start*.

This section includes information on using the PlateSpin® Recon flexible Chargeback feature.

- ♦ [Section 7.1, “Understanding Chargeback,” on page 141](#)
- ♦ [Section 7.2, “Using Raters,” on page 141](#)
- ♦ [Section 7.3, “Managing Raters,” on page 143](#)
- ♦ [Section 7.4, “Chargeback Reporting,” on page 146](#)

7.1 Understanding Chargeback

The PlateSpin Recon flexible chargeback reporting capabilities allow organizations to accurately calculate IT costs based on actual and virtualized resource usage. IT departments can then allocate their costs to different business units and departments.

Even when resources are part of a virtualized pool of computing resources shared across many workloads, PlateSpin Recon not only allows you to track, assign, and report on used resources, but it can also generate reports and deliver them to the appropriate business units.

To set up a chargeback system, first create raters to assign appropriate charges for usage. Next, attach the raters to particular machines to be tracked. Last, generate Chargeback reports to provide the desired information.

NOTE: It is useful to organize machines using Chargeback into their own group. For more information, see [Section 3.4, “Working with Groups,” on page 29](#).

7.2 Using Raters

A rater is a metric used to assign a billable charge to the use of a resource—or part of a resource—over a set amount of time. This converts resource usage data into currency values. There are three types of raters:

- ♦ **Allocation Raters:** These raters can only be attached to virtual machines. They take a cost for a virtual machine server resource and assign an appropriate percentage of that cost to each virtual machine the rater is attached to, based on the usage of that resource.
 - ♦ **Flat Raters:** These raters assign a value to a resource at a fixed rate. For example, a flat rater charge for used disk space could be \$1 per gigabyte of disk space used per month.
 - ♦ **Tiered Raters:** These raters assign values to a resource with different flat rates for different levels of usage. For example, a tiered rater charge for memory used could be \$0 for 0-512 MB and \$.05 per MB for use per day above 512 MB.
- ♦ [Section 7.2.1, “Rater Resources,” on page 142](#)
 - ♦ [Section 7.2.2, “Rater Charges,” on page 142](#)
 - ♦ [Section 7.2.3, “Virtual Server Rater Methods,” on page 143](#)

7.2.1 Rater Resources

Raters calculate charges either according to the maximum use of a resource, or according to the average use during a specific time period. The exception is when they charge against hourly data, where charges are calculated using the average use only.

Table 7-1 *Rater Resources*

Resource	Description
Uptime (Hours)	The charge is based on the number of hours the network entity has been monitored by PlateSpin Recon during the time specified in the report. <i>Measurement and Time Period are dimmed for Flat and Tiered Raters. Method is also dimmed for Allocation Raters.</i>
VC_ProcessorTime (hour)	The hourly average usage is computed and the sum of all the hour averages during the time specified in the report is used for the charge. <i>Measurement and Time Period are dimmed.</i>
Processor (MHz)	Selecting the Hour <i>Time Period</i> forces the use of the Average value, even if Maximum is selected.
Memory Used(MB)	
Network (MB/sec)	
Disk (MB/sec)	
Disk Space Used (GB)	
Disk Writes (MB/sec)	
Total Processor (Count)	Static counter. The charge is made against a static value. <i>Measurement and Time Period are dimmed.</i>
Total Memory (MB)	Static counter. The charge is made against a static value. <i>Measurement and Time Period are dimmed.</i>
Total Disk (GB)	Static counter. The charge is made against a static value. <i>Measurement and Time Period are dimmed.</i>

NOTE: Aside from *VC_ProcessorTime*, described in the preceding table, VC_specific counters are not required nor useful for chargeback, and as a result are not available when creating raters. Use the standard memory counters (*Memory Used* and *Total Memory*) as needed.

7.2.2 Rater Charges

Flat and Tiered raters: Charge represents the amount of cost assigned to a machine for every measure of usage. For example, for every 3 GB of disk space used per week, apply the charge. For Tiered raters, a Charge value must be entered for each tier.

Allocation raters: Charge represents the total cost of a virtual machine server resource. For example, a rater could include a charge of \$1000 per month for virtual machine server disk space. Every virtual machine to which the rater is attached is assigned a percentage of the \$1000, proportional to its usage of the selected resource. For example, based on the above rater, if a virtual machine server contains 3 virtual machines that use 50%, 25%, and 10% of its disk space, respectively (either Virtual Server Usage or Virtual Server Capacity), they are assigned \$500, \$250, and \$100 per month in chargeback costs.

IMPORTANT: When using Chargeback with Solaris Zones, allocation raters calculate utilization charges against the host, not the zone, using the formula:

- ♦ $\text{charge} = \text{unit price} * (\text{zone total resource} / \text{host total resource})$

If the zone and the host have the same number of resources—CPU, disk or memory—then the result is as expected; but if they differ, unexpected results could be seen in the Cost Allocation report.

For example, if the host has 2 CPUs but the zone only has 1 and the charge in the rater for the Processor resource is specified as \$1, then with the above formula the resulting charge would be \$0.50.

If the host has a considerably larger number of resources compared to the zone, especially when looking at memory and disk resources, the resulting charge could be very small. Since the Chargeback Cost Allocation report displays only two decimal places, the conversion could theoretically lose these small values and the charge on the report could appear as zero.

7.2.3 Virtual Server Rater Methods

The Virtual Server Rater Method determines whether the cost assigned to any virtual machine the rater is attached to is calculated based on its percent of usage of the virtual machine server total Virtual Server Usage for that resource, or based on its percent of usage of the virtual machine server total Virtual Server Capacity for that resource. For example, if a virtual machine server has 1000 GB of disk space, but only 800 are actually used, a virtual machine using 200 GB of disk space is using 20% of Virtual Server Capacity, or 25% of Virtual Server Usage. For a Charge of \$1000/month, the virtual machine would be charged \$200 if Virtual Server Capacity is selected, but \$250 if Virtual Server Usage is selected.


7.3 Managing Raters

The Rater Explorer window is used to manage raters. You can add, edit, or delete raters, as well as attach or detach raters to or from discovered machines. The Rater Explorer also allows you to view current raters and their attachments.

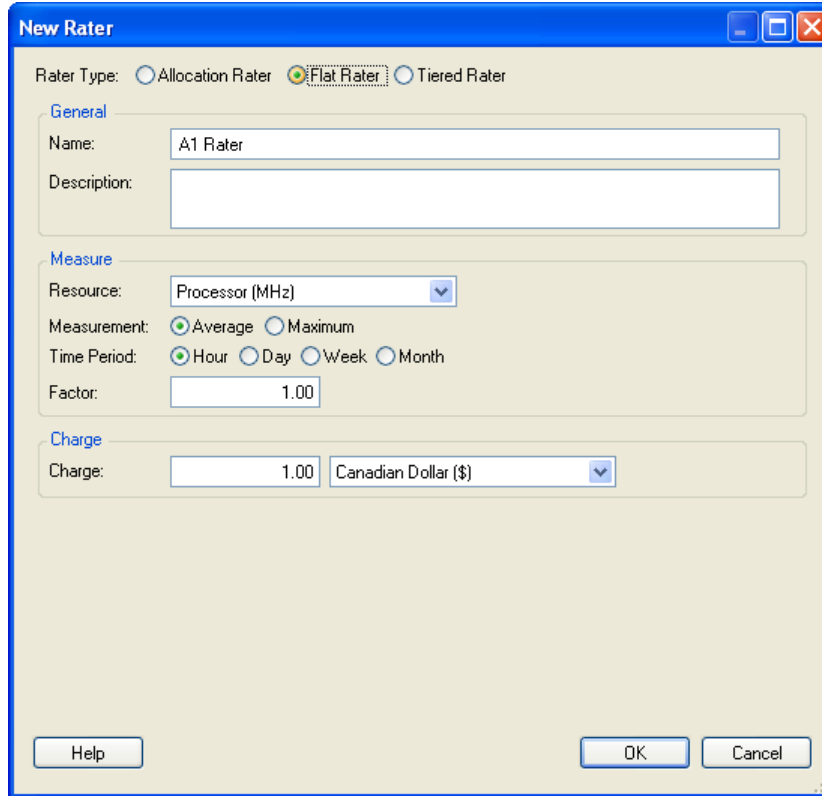
A machine can have multiple raters and a rater can be used with multiple machines. Ratings are performed on a one machine to one rater basis (that is, a rater only considers data from one machine at a time).

- ♦ [Section 7.3.1, “Creating a Rater,” on page 144](#)
- ♦ [Section 7.3.2, “Deleting a Rater,” on page 145](#)
- ♦ [Section 7.3.3, “Editing a Rater,” on page 145](#)
- ♦ [Section 7.3.4, “Attaching Raters to Machines,” on page 145](#)
- ♦ [Section 7.3.5, “Detaching a Rater from a Machine,” on page 146](#)

7.3.1 Creating a Rater

- 1 In the Rater Explorer, click the *New* button .

The New Rater dialog box is displayed..



The "New Rater" dialog box is shown with a blue title bar and standard window controls. It contains three sections: "General", "Measure", and "Charge". In the "General" section, "Rater Type" has three radio buttons: "Allocation Rater", "Flat Rater" (which is selected), and "Tiered Rater". Below this are text fields for "Name" (containing "A1 Rater") and "Description". The "Measure" section has a "Resource" dropdown menu set to "Processor (MHz)", "Measurement" radio buttons for "Average" (selected) and "Maximum", "Time Period" radio buttons for "Hour" (selected), "Day", "Week", and "Month", and a "Factor" text field set to "1.00". The "Charge" section has a "Charge" text field set to "1.00" and a currency dropdown menu set to "Canadian Dollar (\$)". At the bottom are "Help", "OK", and "Cancel" buttons.

- 2 Select a *Rater Type*.

For more information on Rater types, see [Section 7.2, “Using Raters,” on page 141](#).

- 3 Type a *Name*.

- 4 (Optional) Type a *Description*.

- 5 Select a *Resource* from the drop-down list.

These are also known as counters, because their purpose is to count usage of a resource. For more information on resources, see [Section 7.2.1, “Rater Resources,” on page 142](#).

- 6 Select the *Measurement* and *Time Period* to use, if they are not dimmed.

For *Measurement*, selecting *Average* means that the average value over the *Time Period* is used. Selecting *Maximum* means that out of all the averages within the *Time Period*, the maximum average is used. For example, there are 24 averages per day for hourly data and the maximum average out of all the 24 averages is used against the charge value.

- 7 Type a *Factor*.

This is the multiple of the resource units selected that you want to use for the charge. For example, you could charge \$2.00 per 3 GB of disk space used. In this case, the *Factor* would be 3.

- 8 Type a *Charge* monetary value and select a currency type from the drop-down list.

For more information, see [Section 7.2.2, “Rater Charges,” on page 142](#).


- 9 (Conditional) For Allocation raters (virtual machines only), select a *Method* from the drop-down list.

For more information, see [Section 7.2.3, “Virtual Server Rater Methods,” on page 143](#).


- 10 Click *OK*.

The new rater is displayed in the Rater Explorer window under the appropriate rater type.

7.3.2 Deleting a Rater

- 1 In the Rater Explorer window, select the rater you want to delete
- 2 Click the *Delete* button .

7.3.3 Editing a Rater

- 1 In the Rater Explorer window, select the rater you want to edit .
- 2 Click the *Edit* button .

The Edit Rater dialog box is displayed.

- 3 Edit the rater values as desired. Use the steps in [Step 1 on page 144](#) as a guide.

The Rater Type is not editable. To change a rater type, you need to delete the rater and then create a new rater with the same name that uses the desired type.

- 4 Click *OK*.

A confirmation prompt window is displayed.

- 5 Click *OK*.


The rater is removed from the Rater Explorer window.

Deleting a rater that is attached to a machine does not affect the machine, except to exclude it from any chargeback calculations unique to the deleted rater.

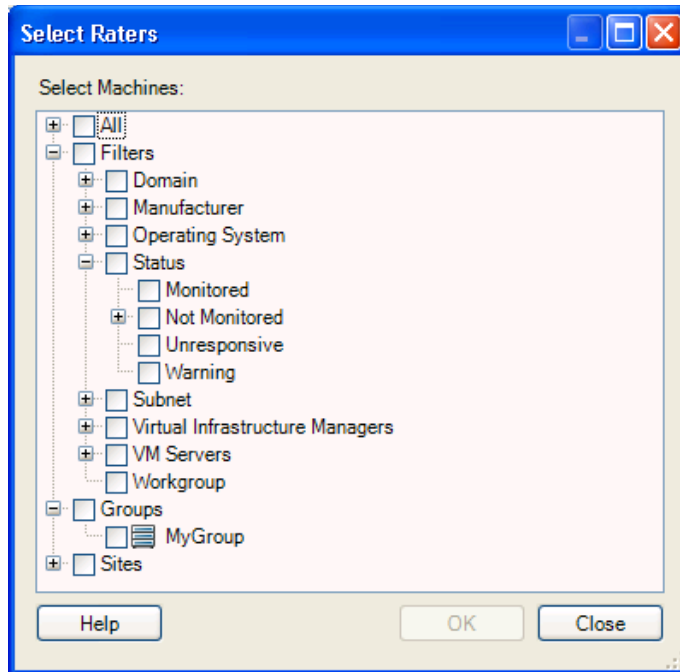
7.3.4 Attaching Raters to Machines

A machine needs to be inventoried and in most cases monitored to enable a rater to access its resources. You then need to attach the rater to the machine to have the machine cost calculated and made available for reporting.

To calculate a charge, PlateSpin Recon must have monitoring data for all counters except Total Processor, Total Memory, and Total Disk, which use inventory data.

- 1 In the Rater Explorer window, select the rater you want to attach to a machine.
- 2 Click the *Attach* button .

The Select Raters dialog box is displayed.



You can also use the Data Center Explorer window to attach a rater to a machine . Right-click the machine you want to add the rater to and click *Chargeback* in the menu. The Raters for <machinename> window is displayed where you can add or remove raters by selecting them and using the *Add* and *Remove* buttons. Adding and removing raters from here is the same as attaching or detaching them from the Rater Explorer window. It does not add or remove the raters from PlateSpin Recon.


- 3 Select the check box next to the machine you want to add the rater to.

You can select more than one machine, if desired. Raters can be attached to groups, sites, or filters by selecting them in the Select Raters dialog box.

- 4 Click *OK*.

The machines you attached the rater to are displayed in the Rater Explorer window beneath that rater.

7.3.5 Detaching a Rater from a Machine

- 1 In the Rater Explorer window, select the machine you want to detach the rater from.
- 2 Click the *Detach* button .

The rater is immediately detached and the machine is no longer displayed beneath the rater in the Rater Explorer window.

7.4 Chargeback Reporting

Use the Cost Allocation report template under the Chargeback node in the Report Explorer window to calculate chargeback values for a particular time period and generate a report that includes chargeback data. In addition to the columns available for inventory reports, Chargeback reports can include raters attached to machines, their types, their calculated costs, and a description.

To generate a report that calculates an IT chargeback cost for a particular business unit, ensure that the *Chargeback - Calculated Cost* column is included in the report. In the *Groups* tab of the Create View dialog box, create a grouping by a column such as *Machine - Department* or *Machine - Billing Code*.

The generated report includes a total chargeback cost for every department, billing code, or any other column used for grouping.

For more information see [Chapter 5, “Generating and Viewing Charts and Reports,” on page 75](#) and [Section 5.3.2, “Chargeback Report Templates,” on page 89](#).

7.4.1 How Calculated Cost is Derived

The *Calculated Cost* in the Cost Allocation report is derived based on the Rater a machine is attached to and the date range specified.

Figure 7-1 Sample Rater

The screenshot shows the 'Edit Rater' dialog box with the following configuration:

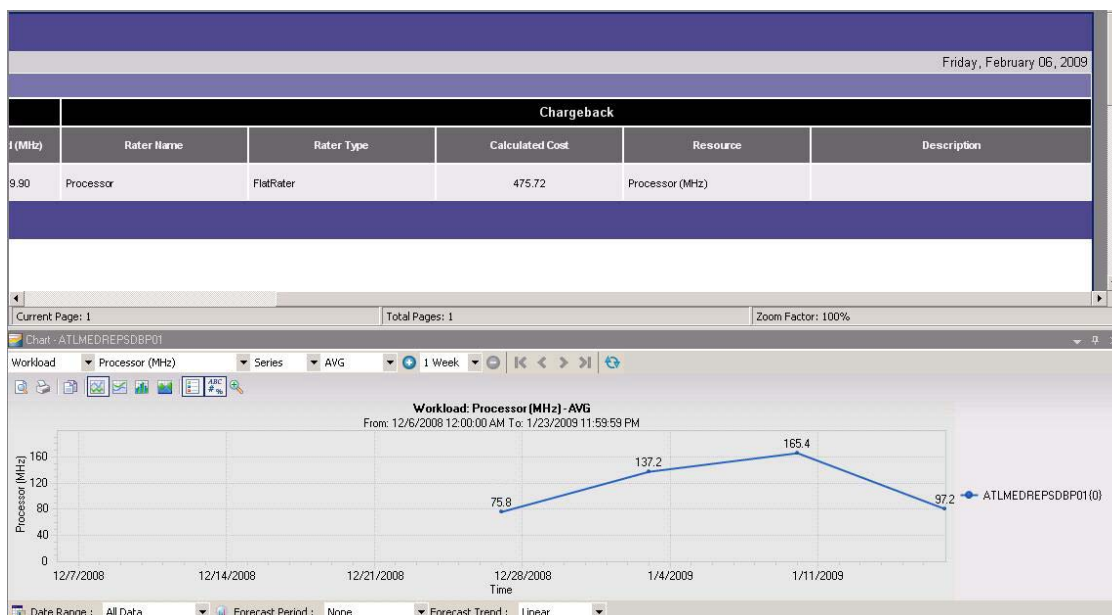
- Rater Type:** ☐ Allocation Rater ☒ Flat Rater ☐ Tiered Rater
- General:**
 - Name: MyRater
 - Description: (empty)
- Measure:**
 - Resource: Processor (MHz)
 - Time Period: ☐ Hour ☐ Day ☒ Week ☐ Month
 - Measurement: ☒ Average ☐ Maximum
 - Factor: 1.00
- Charge:**
 - Charge: 1.00
 - US Dollar (USD)

Buttons at the bottom: Help, OK, Cancel.

A machine attached to the above Rater would be charged \$1.00 for every week of 1 MHz processor use.

When the Cost Allocation report is run for this machine, PlateSpin Recon adds the weekly processor use together for each week in the report and multiplies it by the charge specified in the Rater. The following figure shows a sample processor use chart and the chargeback Cost Allocation report that would result.

Figure 7-2 Sample Chart and Correlating Chargeback Report



The four data points are added together, multiplied by the unit cost, giving you the Calculated Cost. The slight difference between this calculation and what is shown in the report is due to the chart rounding values to a single decimal place. In this case:

- ♦ $(75.8 + 137.2 + 165.4 + 97.2) * 1.00 = 475.6$

The PlateSpin® Recon Virtual Infrastructure Management tools facilitate the administration of Virtual Infrastructure in the data center. These include Virtual Center monitoring and the ability to set up recurrent inventory jobs to have the most up-to-date information about virtual machines being managed by a Virtual Center. VM Sprawl reports can also be generated on a one-time or recurrent basis to provide information such as the age, last modification, and last boot time of the virtual machines in the data center.

- ♦ [Section 8.1, “Virtual Center Inventory and Monitoring,” on page 149](#)
- ♦ [Section 8.2, “Scheduling and Recurrence,” on page 150](#)
- ♦ [Section 8.3, “VM Sprawl Reporting,” on page 152](#)
- ♦ [Section 8.4, “Virtual Machine State Reporting,” on page 152](#)

8.1 Virtual Center Inventory and Monitoring

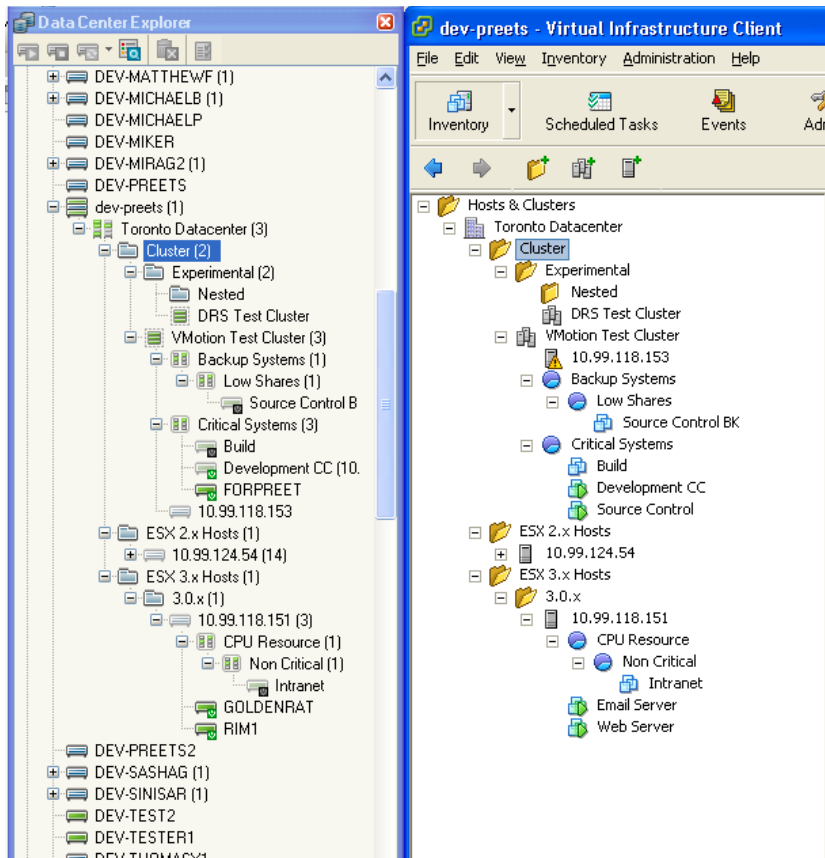
Direct Virtual Center inventory and monitoring enables multiple virtual machines to be more efficiently inventoried from one single point instead of inventorying them individually.

Monitoring through the Virtual Center provides different information than monitoring a virtual machine directly, because it is a picture of the virtual machine resource usage as detected by the Virtual Center and ESX host. This is more accurate than a virtual machine reporting on its own resource utilization. For example, a virtual machine using 90% of its allotted processing capabilities reports its processor usage as 90%. In reality, it might have been allocated only 30% of the host processor capabilities. This difference can only be detected by monitoring through the Virtual Center, which reports what percentage of the host processor a virtual machine is using.

This difference has become more and more useful as data centers place emphasis on managing and optimizing virtual machine usage as opposed to merely consolidating physical servers onto virtual machines.

When a Virtual Center is inventoried, its associated machines are displayed in the Data Center Explorer in a hierarchy that reflects the VMware Virtual Infrastructure Client hierarchy, including all data centers, folders, clusters and pools.

Figure 8-1 Data Center Explorer and VMware Virtual Infrastructure Client Comparison



Inventorying a Virtual Center discovers all the machines that are part of the Virtual Center, so that they can then be inventoried and monitored. When a virtual machine is inventoried through the Virtual Center, it shows up as inventoried both as a node in the Virtual Center, and as an inventoried machine under All in the Data Center Explorer. Either node can subsequently be used to initiate monitoring.

You can right-click a cluster in a Virtual Center to view its properties, which are set in the Virtual Center.

When a virtual machine that is part of a Virtual Center is monitored, certain values, such as Processor Time, are collected both directly from the virtual machine and through the Virtual Center. Values that are being collected through the Virtual Center are prefixed with VC, for example VC_ProcessorTime.

See [Section 3.7.2, “Virtual Center,” on page 33](#) for information about how to inventory a Virtual Center.

8.2 Scheduling and Recurrence

Certain PlateSpin Recon jobs can be scheduled to occur at a later date, or on a recurring basis. These can be viewed through the Scheduled Job Explorer. Jobs can include inventorying, report generation, and snapshot export.

Scheduling is controlled through options in the Export Site Snapshot (Step 3 of 3) and Generate Reports dialog boxes. For more information see [Section 3.13.1, “Exporting a PlateSpin Recon Snapshot,” on page 47](#) and [Section 5.5, “Generating Reports,” on page 96](#).

The same options are available in the Inventory Machines dialog box by clicking *Schedule* on the drop-down menu for the *Inventory* button. The Scheduling Options dialog box is displayed.

Figure 8-2 Scheduling Options

To schedule a job:

- 1 In the Scheduling Options dialog box, Select *Later* to run the job at a specific date and time (either typed manually in mm/dd/yyyy hh:mm AM|PM format, or selected from the pop-up calendar). Continue with [Step 5](#).

or

Select *Recurrence* to run the job on a recurring basis. From the drop-down list, select a previously entered recurrence schedule, or select *<New Recurrence>*.

The Create Recurrence dialog box is displayed.

- 2 Enter a unique *Recurrence Name*.

This name is used to refer to the recurrence schedule in the *Recurrence* drop-down list, so that it can be reused for other jobs.

- 3 In the *Recurrence pattern* area, select a rate of recurrence, and then use options to specify when, and how often jobs are run.
- 4 In the *Range of recurrence* area, specify a start date and time, and then indicate how the recurrence is to extend. Specify the total number of occurrences for the job, or set an end date.
By default, the recurrence extends indefinitely and can only be stopped by right-clicking the job in the Scheduled Job Explorer and clicking either *Delete* or *Edit Recurrence* on the menu.
- 5 Click *OK*.

8.3 VM Sprawl Reporting

The VM Sprawl report provides information about all virtual machines found by PlateSpin Recon, along with information about their servers.

NOTE: If a VM is under a cluster and the parent host ESX machine is not inventoried by PlateSpin Recon, the VM is displayed under the Cluster group of this report.

You can use this report to keep track of the number and location of virtual machines resident in the data center. VM Sprawl report information is obtained from the virtual machines' host servers.

Citrix XenServer pools can also be inventoried by PlateSpin Recon, and these pools are treated similarly to a Virtual Center. The VM Sprawl report can also be used to report on these pools.

The VM Sprawl Report template is located in the Inventory Reports section of the Reports Explorer. For more information about Inventory reports, see [Section 5.2.4, "Working with Report Templates," on page 83](#) and ["Deleting a Report Template" on page 84](#).

8.4 Virtual Machine State Reporting

The Virtual Machine State report shows only the machines displayed under the VM Servers filter. It is found under the Inventory Reports section of the Reports Explorer.

Technical Reference

A

This section contains technical information to help you customize PlateSpin® Recon.

- ♦ [Section A.1, “Features,” on page 153](#)
- ♦ [Section A.2, “How PlateSpin Recon Collects Data,” on page 153](#)
- ♦ [Section A.3, “Data Collection,” on page 155](#)
- ♦ [Section A.4, “Charts and Formulas,” on page 156](#)
- ♦ [Section A.5, “Calculations,” on page 160](#)
- ♦ [Section A.6, “CPU Normalization,” on page 162](#)
- ♦ [Section A.7, “PlateSpin Recon VM CPU Inventory,” on page 163](#)
- ♦ [Section A.8, “Linux/Solaris Inventory and Monitoring Commands,” on page 164](#)

This section is for system administrators and virtualization architects who require advanced knowledge of PlateSpin Recon.

It is assumed that readers of this document are familiar with:

- ♦ Administration of Windows and Linux operating systems
- ♦ Administration of VMware ESX, VMware Server, and Microsoft Virtual servers

A.1 Features

- ♦ Hardware and software asset inventory
- ♦ Secure on-site data collection and analysis
- ♦ Windows and Linux support
- ♦ Management reports
- ♦ Integration with Microsoft Operations Manager 2005 and the remote PlateSpin Recon server
- ♦ Advanced project and scenario modeling
- ♦ Integration with the PlateSpin Portability Suite

A.2 How PlateSpin Recon Collects Data

PlateSpin Recon has three sequential stages to its data collection.

- ♦ [Section A.2.1, “Discovery,” on page 153](#)
- ♦ [Section A.2.2, “Inventory,” on page 154](#)
- ♦ [Section A.2.3, “Monitoring,” on page 155](#)

A.2.1 Discovery

Domain Discovery: PlateSpin Recon uses Windows Active Directory* via LDAP to scan the network for a list of the machines on the specified domain. By default, this includes only online machines, but there is an option to include offline machines as well.

An Organizational Unit (OU) filter can also be specified, narrowing the area of the domain that PlateSpin Recon will poll during discovery. An Organization Unit is a container within a domain where computers can reside for segmentation. For example, if your domain has OU containers setup for each department, you can tell PlateSpin Recon to just look for machines within a specific department within the domain.

PlateSpin Recon only uses OU filters during discovery. Machines discovered in this way are unaffected during inventory and monitoring should machines be moved out of their previous OU containers. For more information on Organization Units and if they are in use in your domains, check with your System Administrator.

Subnet, IP Range Scan: For each machine in the subnet or IP range, PlateSpin Recon pings the machine. If it replies, it is considered a discovered machine.

Another option is to port scan through TCP, UDP or both. PlateSpin Recon tries to connect to ports and records which ports are being used. This option must be used with caution because network security might consider this an attack.

A.2.2 Inventory

- ♦ “Linux, Solaris and ESX 2.x” on page 154
- ♦ “ESX 3.x and Virtual Center” on page 154
- ♦ “Microsoft Windows Inventory” on page 154

Linux, Solaris and ESX 2.x

- ♦ PlateSpin Recon sends the `getplatform` script, which returns the architecture and `glibc` version of the machine being inventoried.
- ♦ Based on `getplatform`, PlateSpin Recon uses the SCP protocol to transfer a platform-specific inventory binary and libraries to the `/tmp` directory of the machine being inventoried.
- ♦ Over `ssh`, PlateSpin Recon executes the binary, streaming the command file over `stdin`.
- ♦ Logs and progress files are streamed back from the inventoried machine to the PlateSpin Recon Server using `stderr` while the machine XML is streamed over `stdout`.

ESX 3.x and Virtual Center

- ♦ PlateSpin Recon runs the executable locally on the PlateSpin Recon Server.
- ♦ The executable accesses ESX 3 or Virtual Center Webservices, which provide the necessary inventory data.

Microsoft Windows Inventory

- ♦ PlateSpin Recon copies an executable onto the machine being inventoried to the a directory within `ADMIN$`. PlateSpin Recon runs the executable through WMI or a remote service. If you are inventorying a Windows NT* machine, it is important to make sure that WMI is installed. By default PlateSpin Recon uses WMI, but it is possible to configure it to use the remote service by using the configuration options in the *Tools* Menu.

A.2.3 Monitoring

- ♦ “Linux, Solaris and ESX 2.x” on page 155
- ♦ “ESX 3.x and Virtual Center” on page 155
- ♦ “Microsoft Windows” on page 155

Linux, Solaris and ESX 2.x

- ♦ PlateSpin Recon sends a script (`lininfo.sh`, `solinfo.sh`, or `esxinfo.sh`) to the machine being inventoried.
- ♦ The script is run through `ssh`.
- ♦ The `ssh` server must be enabled for monitoring to function.
- ♦ Logs are streamed back to the PlateSpin Recon Server over `stderr`.
- ♦ Performance data is streamed back over `stdout`.

ESX 3.x and Virtual Center

- ♦ PlateSpin Recon calls ESX 3 or Virtual Center Webservices, which provide the necessary performance data.

Microsoft Windows


- ♦ PlateSpin Recon uses the Windows Performance Counter API to retrieve performance data. It does not use WMI.
- ♦ The RemoteRegistry Service must be enabled for Windows monitoring to function.

A.3 Data Collection

By default, every 5 minutes, PlateSpin Recon takes three samples, three seconds apart, from all counters. The three samples are averaged together to produce the five-minute data that is stored in the database.

These values can be modified by clicking the *Tools* menu, then holding down the Ctrl key while selecting *Options* from the menu. On the *Server* tab, expand the Monitoring Service section:


Figure A-1 Options Window



Burst interval: The interval between different burst samples

Burst samples: The number of samples collected at a burst interval before an average is calculated

Sample interval: The interval between different samples



A.4 Charts and Formulas

- ♦ Section A.4.1, “Chart Data,” on page 157
- ♦ Section A.4.2, “Utilization,” on page 158
- ♦ Section A.4.3, “Workload,” on page 158

A.4.1 Chart Data

The five-minute data collected as described in [Section A.3, “Data Collection,” on page 155](#) is stored in the database for thirty days. The one-hour data is a summary of all the five-minute data collected during a particular hour. It is stored permanently in the database. All other chart data is calculated from the five-minute and one-hour data.

The twenty-four hour profile chart provides a value for each hour in a twenty-four hour period. This value is the largest average value calculated for the particular hour over the selected date range. This produces a utilization profile for the hourCounter.

Counter values are named for the Windows Performance Monitor counters used to retrieve them from Windows systems. The name consists of the Performance Monitor category joined by an underscore to the counter name. For example, Memory/PagesPerSecond becomes Memory_PagesPerSecond.

For Windows systems, these values are taken directly from Performance Monitor, but for other platforms, scripts are used to calculate the values.

Table A-1 Platform Counters

Name	Platforms
LogicalDisk_FreeMegabytes	Windows
LogicalDisk_PercentFreeSpace	Windows
Memory_AvailableBytes	Windows, Windows NT, Solaris, Linux, ESX 2.5.x, ESX 3.x
Memory_PagesPerSecond	Windows, Windows NT
Memory_PercentAvailableMemory	ESX 3.x
NetworkInterface_BytesTotalPerSecond	Windows, Solaris, Linux, ESX 2.5.x, ESX 3.x
NetworkInterface_PacketsPerSecond	Windows, Solaris, Linux, ESX 2.5.x, ESX 3.x
PhysicalDisk_CurrentDiskQueueLength	Windows, Windows NT, Solaris, Linux, ESX 2.5.x, ESX 3.x
PhysicalDisk_DiskBytesPerSecond	Windows, Windows NT, Solaris, Linux, ESX 2.5.x, ESX 3.x
PhysicalDisk_DiskTransfersPerSecond	Windows, Windows NT, Linux, ESX 2.5.x
PhysicalDisk_DiskWriteBytesPerSecond	Windows
PhysicalDisk_PercentIdleTime	Windows
Processor_PercentProcessorTime	Windows, Windows NT, Solaris, Linux, ESX 2.5.x, ESX 3.x
System_ProcessorQueueLength	Windows, Windows NT, Solaris, Linux, ESX 2.5.x
VC_Processor	All
VC_ProcessorTime	All
VC_MemoryPercentAvailable	All

Name	Platforms
VC_MemoryAvailableBytes	All
VC_NetBytesTotalPerSecond	All
VC_NetPacketsPerSecond	All
VC_DiskBytesPerSecond	All

A.4.2 Utilization

Utilization formulas represent the percentage of the available resources being used by a given workload. The calculations are based on counter and inventory data.

Table A-2 Utilization Formulas

Name	Formula	Platforms
Disk IO (%)	MAX(0,(100-PERF("PhysicalDisk_PercentIdleTime")))	Windows
Disk Space Free (%)	100*(INV("FreeDiskSpace")/INV("TotalDiskSpace"))	Windows,Solaris, Linux, ESX 2.5.x, ESX 3.x
Disk Space Used (%)	100*(INV("TotalDiskSpace")-INV("FreeDiskSpace")/INV("TotalDiskSpace"))	Windows, Solaris,Linux, ESX 2.5.x, ESX 3.x
Memory Free (%)	100*(PERF("Memory_AvailableBytes")/INV("TotalMemory"))	Windows, Solaris,Linux, ESX 2.5.x, ESX 3.x
Memory Used (%)	100*(INV("TotalMemory")-PERF("Memory_AvailableBytes")/INV("TotalMemory"))	Windows,Solaris, Linux, ESX 2.5.x, ESX 3.x
Processor (%)	PERF("Processor_PercentProcessorTime")	Windows, Solaris, Linux, ESX 2.5.x, ESX 3.x
VC_MemoryFree (%)	100*(PERF("Memory_AvailableBytes")/INV("TotalMemory"))	All
VC_MemoryUsed (%)	100*(INV("TotalMemory")-PERF("Memory_AvailableBytes")/INV("TotalMemory"))	All
VC_ProcessorTime (Hour)	PERF("VC_ProcessorTime")	All

A.4.3 Workload

Workload formulas measure workload size; that is, the amount of processing power and raw resources being used for a workload, independent of the total hardware and resources available.

- ♦ PERF: Dynamic performance data
- ♦ INV:Static inventory data

Table A-3 *Workload Formulas*

Name	Formula	Platforms
Disk (MB/sec)	$(\text{PERF}(\text{"PhysicalDisk_DiskBytesPerSecond"}))/(1024*1024)$	All
Disk (Transfers/sec)	$\text{PERF}(\text{"PhysicalDisk_DiskTransfersPerSecond"})$	Windows, Windows NT, Solaris, Linux, ESX 2.5.x
Disk Queue Length	$\text{PERF}(\text{"PhysicalDisk_CurrentDiskQueueLength"})$	Windows, Windows NT, Solaris, Linux, ESX 2.5.x
Disk Space Free (GB)	$\text{INV}(\text{"FreeDiskSpace"})/(1024*1024*1024)$	All
Disk Space Used (GB)	$(\text{INV}(\text{"TotalDiskSpace"}) - \text{INV}(\text{"FreeDiskSpace"}))/(1024*1024*1024)$	All
Disk Writes (MB/sec)	$(\text{PERF}(\text{"PhysicalDisk_DiskWriteBytesPerSecond"}))/(1024*1024)$	Windows
Memory (Pages/sec)	$\text{PERF}(\text{"Memory_PagesPerSecond"})$	Windows, Windows NT
Memory Free (MB)	$(\text{PERF}(\text{"Memory_AvailableBytes"}))/(1024*1024)$	All
Memory Used (MB)	$(\text{INV}(\text{"TotalMemory"}) - \text{PERF}(\text{"Memory_AvailableBytes"}))/(1024*1024)$	All
Network (MB/sec)	$(\text{PERF}(\text{"NetworkInterface_BytesTotalPerSecond"}))/(1024*1024)$	Windows, Solaris, Linux, ESX 2.5.x, ESX 3.x
Network (Packets/sec)	$\text{PERF}(\text{"NetworkInterface_PacketsPerSecond"})$	Windows, Solaris, Linux, ESX 2.5.x, ESX 3.x
Processor (MHz)	$\text{PERF}(\text{"Processor_PercentProcessorTime"}) * \text{INV}(\text{"ProcessorCapacity"}) / 100$	Windows, Windows NT, Linux, ESX 2.5.x, ESX 3.x
Processor Queue Length	$\text{PERF}(\text{"System_ProcessorQueueLength"})$	Windows, Windows NT, Solaris, Linux, ESX 2.5.x
VC_Disk (MB/sec)	$(\text{PERF}(\text{"PhysicalDisk_DiskBytesPerSecond"}))/(1024*1024)$	All
VC_MemoryFree (MB)	$(\text{PERF}(\text{"Memory_AvailableBytes"}))/(1024*1024)$	All
VC_MemoryUsed (MB)	$(\text{INV}(\text{"TotalMemory"}) - \text{PERF}(\text{"Memory_AvailableBytes"}))/(1024*1024)$	All
VC_Network (MB/sec)	$(\text{PERF}(\text{"NetworkInterface_BytesTotalPerSecond"}))/(1024*1024)$	All

Name	Formula	Platforms
VC_Network (Packets/sec)	PERF("NetworkInterface_PacketsPerSecond")	All
VC_Processor (MHz)	PERF("Processor_PercentProcessorTime")*INV("ProcessorCapacity")/100	All

A.5 Calculations

- ♦ [Section A.5.1, “24-Hour Profile,” on page 160](#)
- ♦ [Section A.5.2, “Dual-Core and Multi-Core Machines,” on page 161](#)
- ♦ [Section A.5.3, “Hyper-Threading,” on page 161](#)
- ♦ [Section A.5.4, “Effective Speed,” on page 162](#)
- ♦ [Section A.5.5, “Power and Cooling,” on page 162](#)

A.5.1 24-Hour Profile

The 24-hour profile consists of 24 values that correspond to each hour of the day (0 to 23). Each value is calculated by taking the peak value for that specific hour over the number of monitoring days.

Example

Consider a system monitored for 3 days. The following chart presents the data set for its Processor workload, and the calculated 24-hour profile:

Table A-4 *Processor Workload Data Set*

Hour	Day 1	Day 2	Day 3	24-Hour Profile
0	200	300	500	500
1	300	300	250	300
2	300	400	450	450
3	700	250	350	750
4	500	500	500	500
5	600	400	500	600
6	500	550	750	750
7	400	600	800	800
8	500	300	250	500
9	450	350	640	640
10	200	240	340	340
11	300	600	+60	650

Hour	Day 1	Day 2	Day 3	24-Hour Profile
12	300	-	-	300
13	550	340	850	850
14	440	500	600	600
15	800	540	600	800
16	750	260	430	750
17	500	640	450	640
18	600	200	350	600
19	300	250	650	650
20	200	300	200	300
21	250	350	400	400
22	600	250	350	600
23	500	650	750	750

A.5.2 Dual-Core and Multi-Core Machines

PlateSpin Recon considers and treats dual-core or multi-core machine as two-processor or x-processor machines.

A.5.3 Hyper-Threading

The effect of hyper-threading is controlled by two configuration parameters: Aggregate Instance Function (SUM, AVG) and Hyper-threading Normalization Factor (1 – 2). The default value for aggregate instance function is AVG, and the normalization factor is set to 1.25. These values can be configured in the advanced settings under *Tools > Options*.

PlateSpin Recon collects each individual processor instance's utilization values at a given time and then calculates the processor utilization based on the aggregate function. If the aggregate function is SUM, then the physical processor utilization is the sum of all logical processor instances' utilization. If the aggregate function is AVG, then the physical processor utilization would be average of all logical processor instances' utilization.

Example

Consider a system with single-core processor with hyper-threading enabled at p MHz speed. For a given period of time, PowerRecon calculates the percentile utilization values for 24H profile values of instance X and Y as follows.

$\{X\} = \{X1, X2, Xi, \dots, X24\}$, Xi is the utilization of the processor instance X at hour i.

$\{Y\} = \{Y1, Y2, Yi, \dots, Y24\}$, Yi is the utilization of the processor instance Y at hour i.

Table A-5 Aggregate Function Calculations

(Aggregate Function, Normalization Factor)	(SUM, f)	(AVG, f)
Reports: Peak %	$\text{Max} \{ X + Y \}$	$\text{Max} \{ (X + Y) / 2 \}$
Reports: Peak (MHz)	$\text{Max} \{ X + Y \} * (1/100) * p * f$	$\text{Max} \{ (X+Y)/2 \} * (1/100) * p * f$
CPM: Workload calculation	$\{X+Y\} * (1/100) * p * f$	$\{(X+Y)/2\} * (1/100) * p * f$

A.5.4 Effective Speed

PlateSpin Recon takes both number of cores and hyper-threading into account when calculating the effective speed of a machine.

Example:

Server template specification: 2 x 3000 GHz, dual core, and hyper-threading enabled

Number of processors: $P = 2 * 2 = 4$

Processor capacity: $C = 3000$ GHz

Hyper-threading normalization factor: $n = 1.25$ (default value)

Scaling factor: $A(p) = P / \{ 1 + \alpha (P-1) \}$, where $\alpha = 0.05$

$A(p) = 4 / \{ 1 + 0.05 * 3 \} = 3.4782$

Effective speed = $A(p) * C * n = 3.4782 * 3000 * 1.25 = 13043$ MHz

A.5.5 Power and Cooling

Energy Use (kWh/yr) = Power Usage (W) x Server Derating Factor x Number of hours in a year
(8760) / 1000

Heat Dissipation (BTU/yr) = Energy Use (kWh/yr) x 3.41

Annual Energy Cost = Energy Cost per kWh x Energy Use (kWh/yr)

A.6 CPU Normalization

CPU normalization allows the actual speed of a CPU to be considered for operations such as consolidation planning, instead of relying on the CPU clock frequency, which can be misleading. Different CPU manufacturers and models provide a wide array of performance results for CPUs with the same clock frequency. For example, there is a considerable difference between a Core 2 and a Pentium IV CPU with a clock frequency of 2.0 GHz. CPU normalization uses the SPEC data industry standard benchmarks to give real-world performance values. This is an important factor in realistic consolidation planning. See [Section A.6.1, “SPEC Data,” on page 163](#) for more information.

In consolidation planning, when a scenario is created, the user is given an option to select between clock frequency and normalized speed as the unit of measurement for processor utilization and capacity. The consolidation planning algorithms use the appropriate unit of measurement when assigning workloads to containers, based on these selections. For more information see [Section 6.3, “Consolidation Scenarios,” on page 120](#).

NOTE: The data required to perform consolidation planning based on normalized speed is actually produced when a project is created. Projects from a pre-3.5 version of PlateSpin Recon, even if upgraded, cannot use normalized speed in consolidation planning (the option is disabled in the scenario creation wizard). To use normalization speeds with these projects, they must be re-created from scratch.

- ♦ [Section A.6.1, “SPEC Data,” on page 163](#)
- ♦ [Section A.6.2, “Normalization Speed,” on page 163](#)

A.6.1 SPEC Data

PlateSpin Recon CPU normalized speed makes use of SPECint* and SPECfp* (98, 2000, 2006) CPU benchmarking data. The PlateSpin Recon CPU Normalization rating for specific CPUs are an average of data available from <http://www.spec.org> (<http://www.spec.org>) for those CPUs as of June 2008. For the latest SPECfp* benchmark results, visit <http://www.spec.org> (<http://www.spec.org>).

A.6.2 Normalization Speed

You can change default normalization speed for inventoried machines from the Custom Fields dialog box. You can also view normalization speed in server templates. Although you can't directly change these values, you can change other values and see how the changes affect the normalization speed. For more information, see [Section 6.2.2, “Server Templates,” on page 114](#).

To view or edit normalization speed for an inventoried machine:

- 1 In the Data Center Explorer, right-click a machine and select *Properties*.
The Properties dialog box for that machine is displayed.
- 2 Click the *Custom Fields* tab.
In the top section you can see the default *Normalized Speed* for the machine.
- 3 Click the *Normalized Speed* value and edit as desired.
To reset the value to its default, click *Reset Reference Values*.
- 4 Click *OK*.

CPU Normalization values can also be used as column values in relevant reports, although the columns are not shown by default. To display normalization-oriented columns, see [Section 5.3, “Editing Report Templates,” on page 84](#) and [Section 5.4.1, “Creating and Editing Report Views,” on page 92](#).

A.7 PlateSpin Recon VM CPU Inventory

When PlateSpin Recon inventories Virtual Machines with one or more CPUs allocated to them, the cores and CPUs are displayed as follows:

Cores: Regardless of the VM guest machine configuration, PlateSpin Recon shows the actual number of cores from the hardware host. This may contradict with third-party tools, such as CPUID.

CPUs: On ESX hosts, there are three possible virtual CPU configurations (1,2 and 4) for VMs. The PlateSpin Recon inventory of the CPUs depends on the VM configuration and the cores on the host machine. ESX allocates as many cores as needed to accommodate the user requirement. PlateSpin Recon reports the minimum number of CPUs required to provide the indicated number of virtual CPUs.

The following table shows some example configurations and what PlateSpin Recon reports for each.

Table A-6 *Sample Virtual CPU Configurations and What PlateSpin Recon Reports*

Host Cores	Guest Virtual CPUs	PlateSpin Recon Reported Cores	PlateSpin Recon Reported CPUs
1	1	1	1
	2	1	2
	4	1	4
2	1	2	1
	2	2	1
	4	2	2
4	1	4	1
	2	4	1
	4	4	1

A.8 Linux/Solaris Inventory and Monitoring Commands

This section contains lists of commands used by PlateSpin Recon during inventory and monitoring.

- ♦ [Section A.8.1, “Inventorying Solaris,” on page 164](#)
- ♦ [Section A.8.2, “Inventorying LINUX,” on page 165](#)
- ♦ [Section A.8.3, “Monitoring,” on page 166](#)
- ♦ [Section A.8.4, “Text Processing,” on page 166](#)

A.8.1 Inventorying Solaris

Table A-7 *Inventorying Commands*

Command	Description
kstat	#Get network connection speed.
chmod	#Change file access permissions for the inventory script that is placed on the machine being inventoried.

Command	Description
rm	#Remove the inventory script that was placed on the machine.
prtdiag	#Get CPU cache information and machine model for Sparc.
psrinfo	#Get processor information such as family, model, step, etc., which are related to x86 mainly.
isainfo	#Get CPU instruction set information such as sse, sse2, mmx, etc., which are related to x86 mainly.
prtconf	#Get configured device information, which is related to L2 cache information of x86, network adapter resolving, and ACPI support.
ifconfig	#Get network related information such as IP address, DHCP, net mask, etc.
ndd	#Get network connection speed.
uname	#Determine kernel version.
mount	#Get boot loader.
pkginfo	#Get installed programs.
ps	#Get process information.
who	#Get run level.
hostname	#Get host name.

A.8.2 Inventorying LINUX

Table A-8 *Inventorying Commands*

Command	Description
chkconfig	#Get service(daemon) information.
df	#Get volume information.
mount	#Get volume information.
e2label	#Get disk label.
route	#Get gateway.
ifconfig	#Get network related information such as IP address, net address, and so on.
egrep	#Get DHCP information.
getcfg (SUSE*)	#Get DHCP for SUSE.
rpm	#Get installed programs.
lspci	#Get PCI device information such as vendor name, device model in human readable format.

Command	Description
iwconfig	#Get wireless network adapter NIC list.
iwlist	#Get wireless bit rate.
ethtool	#Get speed and link status for network connection.
dmesg	#Get real MAC address.
kudzu (Redhat*)	#Get network device information for Redhat Linux.
hwnfo (SUSE)	#Get network adapter information for SUSE Linux.
fdisk	#Determine if a partition is active or not.
runlevel	#Get current run level.
hostname	#Get host name and DNS domain name.

A.8.3 Monitoring

Table A-9 *Monitoring Commands*

Command	Description
psrinfo (SunOS*)	#count processor number.
mpstat (SunOS)	#Get processor time per cpu instance.
vmstat (SunOS)	#Get processor queue length and memory available.
ifconfig	#Get NIC list.
sed	#Edit input stream.
tail	#Edit input stream.
netstat (SunOS)	#Get network package transfer rate.
kstat (SunOS)	#Get NIC speed.
prtconf (SunOS)	#Get module and instance binding for NIC.
iostat (SunOS)	#Get disk io information such as queue and speed.
expr	#evaluate express.

A.8.4 Text Processing

grep, cat, cut, awk, ls