

Novell Archive and Version Services

2.1 Administration Guide

Open Enterprise Server 11

July 10, 2012

Novell.

Legal Notices

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

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

Any products or technical information provided under this Agreement may be subject to U.S. export controls and the trade laws of other countries. You agree to comply with all export control regulations and to obtain any required licenses or classification to export, re-export or import deliverables. You agree not to export or re-export to entities on the current U.S. export exclusion lists or to any embargoed or terrorist countries as specified in the U.S. export laws. You agree to not use deliverables for prohibited nuclear, missile, or chemical biological weaponry end uses. 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 ©2007-2012 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.
1800 South Novell Place
Provo, UT 84606
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).

Contents

About This Guide	7
1 Overview of Archive and Version Services	9
1.1 Individual File Losses Impact Business	9
1.2 Benefits of Archive and Version Services	10
1.3 Key Concepts of Archive and Version Services	11
1.3.1 The Archive Server	11
1.3.2 The Archive Database	12
1.3.3 NSS Volumes	13
1.3.4 Versioning Jobs	13
1.3.5 Job Schedules	14
1.3.6 File Versions	14
1.3.7 Delete Policy	14
1.3.8 Archive Versioning Plug-In for Novell iManager	14
1.3.9 Versioning Policy	15
1.3.10 NSS File Version Utility	15
1.3.11 NetStorage Archive Function	15
1.4 Scenarios for Using Archive and Version Services	15
1.4.1 Cross-Platform Work Environments	15
1.4.2 Group Collaboration	16
1.4.3 File and Directory Name Changes	16
1.4.4 Selective File Versioning	16
1.5 What's Next	17
2 What's New or Changed in Archive and Version Services	19
3 Planning for Archive and Version Services	21
3.1 Assessing Your Versioning Needs	21
3.2 Designing the Archive and Version Services Topology	23
3.3 Understanding Archive Server Properties	24
3.4 Understanding Job Properties	25
3.4.1 Job Information	25
3.4.2 Source Server Information	26
3.4.3 Run Schedule	27
3.4.4 Delete Policy	28
3.4.5 Filter	29
3.5 What's Next	32
4 Prerequisites and Guidelines	33
4.1 Network Architecture Prerequisites and Guidelines	33
4.2 Server Prerequisites and Guidelines	33
4.3 Storage Media Prerequisites and Guidelines	34
4.4 Fault Tolerance Guidelines	35
4.4.1 Multiple Connection Channels for Storage Devices	35
4.4.2 Software RAID Devices	35
4.4.3 Server Clusters	35
4.5 Job Guidelines	35

4.6	Schedule Guidelines	36
4.7	Prerequisites for Using the NSS File Version Utility	36
4.8	Prerequisites for Using the NetStorage Archive Function	37
4.9	Prerequisites and Guidelines for Retrieving File Versions	37
4.10	Guidelines for Availability of File Versions	38
4.11	What's Next	39
5	Setting Up Archive and Version Services	41
5.1	Before you begin	41
5.2	Installing and Configuring an Archive Server	41
5.2.1	Unsupported Service Combinations	42
5.2.2	YaST Install	42
5.2.3	Creating an NSS Pool	44
5.2.4	Creating NSS Volumes	44
5.2.5	Configuring Jobs using iManager	45
5.2.6	Restoring Versions	45
5.3	Upgrade	46
5.3.1	Before You Upgrade	46
5.3.2	Post-Upgrade	46
5.3.3	Caveats on Upgrading from OES 2 to OES 11	47
5.4	Reinstall	48
5.5	What's Next	48
6	Managing the Archive Server	49
6.1	Starting ArkManager	49
6.2	Stopping ArkManager	49
6.3	Updating the Archive Server Configuration	50
6.3.1	Updating the Archive Server Configuration with Modified eDirectory Credentials	50
6.3.2	Updating the Archive Server Configuration with Modified PostgreSQL Credentials	50
6.4	Vacuuming the Archive Database	51
6.5	Backing Up the Archive Database	51
6.6	Backing Up the Archive Data	51
7	Configuring Jobs in iManager	53
7.1	Accessing the Archive Versioning Plug-In in iManager	53
7.2	Configuring Archive Volume	54
7.3	Configuring or Viewing Archive Server Properties	54
7.3.1	Accessing the Server Properties Page	54
7.3.2	Setting Archive Server Properties	55
7.3.3	Viewing Authentication Properties	55
7.4	Configuring Default Job Settings	56
7.4.1	Accessing the Default Job Settings Page	56
7.4.2	Setting Default Job Information	56
7.4.3	Setting Default Source Server Information	56
7.4.4	Setting Default Run Schedule Information	57
7.4.5	Setting Default Delete Policy Information	58
7.4.6	Applying Default Job Settings	59
7.5	Configuring Job Properties	60
7.5.1	Accessing the Job Settings Page	60
7.5.2	Setting Job Information	61
7.5.3	Setting Source Server Information	61
7.5.4	Setting Run Schedule Information	62
7.5.5	Setting Delete Policy Information	63

7.6	Creating a Job	65
7.7	Editing a Job's Properties	67
7.7.1	Correcting Missing or Invalid Information	67
7.7.2	Modifying Job Settings	68
7.7.3	Applying Modified Job Settings	68
7.8	What's Next	69
8	Managing Jobs	71
8.1	Viewing a Jobs Report	71
8.1.1	Generating the Jobs Report	71
8.1.2	Sorting the Jobs Report by Column	72
8.1.3	Setting the Refresh Rate for the Jobs Report	72
8.1.4	Understanding the Report Content	72
8.2	Viewing a Job's Details	73
8.3	Starting or Scheduling a Job	74
8.4	Stopping a Job	75
8.5	Viewing the Archive Log	76
8.5.1	Understanding Log Entries	76
8.5.2	Viewing Messages in the Archive Log	76
8.5.3	Viewing Messages in the Server Log File	78
8.6	Filtering Messages in the Archive Log	78
8.6.1	Filtering Log Entries by Severity Type	78
8.6.2	Filtering Log Entries by Job	78
8.6.3	Setting the Display Details	79
8.6.4	Setting the Date and Time Range	79
8.6.5	Resetting the Log Filters	79
8.7	Cleaning Up the Job's User List	79
8.8	Deleting File Versions	80
8.9	Deleting a Job	81
8.10	Viewing a Deleted Jobs Report	81
8.11	Salvaging a Deleted Job	82
8.12	Purging a Deleted Job	83
9	Running Archive and Version Services in a Virtualized Environment	85
10	Configuring Archive and Version Service for Novell Cluster Services	87
10.1	Benefits of Configuring Archive and Version Service for High Availability	87
10.2	Prerequisites	87
10.3	Installing and Configuring Archive and Version Services	88
10.3.1	Creating NSS Shared Disk Partition	88
10.3.2	Creating Linux POSIX File System Shared Disk Partitions	89
10.3.3	Configuring Archive and Version Service on the Shared Disk	90
10.4	Novell Cluster Services Configuration and Setup	90
10.4.1	Archive and Version Resource Template	90
10.4.2	Creating an Archive Service Cluster Resource	91
10.4.3	Archive Load, Unload and Monitor Scripts	92
10.5	Configuring Jobs using iManager	95
10.6	Managing Archive Resource	95

11 Migrating Novell Archive and Version Services to OES 11	97
12 Security Considerations for Archive and Version Services	99
12.1 NSS Encrypted Volumes	99
12.2 Secure Data Transfer	99
12.3 eDirectory and PostgreSQL Administrator Passwords	99
12.4 User Authentication	100
A XML Elements and Attributes for ArkConfig	101
B Sample Configuration Files	117
B.1 arkConfig.xml - Sample Configuration file	117
C Troubleshooting	119
C.1 Archive Server	119
C.2 Migration Issues	123
C.3 NSS File Version Utility	123

About This Guide

This guide describes how to configure and manage a Novell Archive and Version Services 2.1 for OES 11 server to archive multiple interval-based versions of files for convenient access and retrieval by users. It is divided into the following sections:

- ♦ Chapter 1, “Overview of Archive and Version Services,” on page 9
- ♦ Chapter 2, “What’s New or Changed in Archive and Version Services,” on page 19
- ♦ Chapter 3, “Planning for Archive and Version Services,” on page 21
- ♦ Chapter 4, “Prerequisites and Guidelines,” on page 33
- ♦ Chapter 5, “Setting Up Archive and Version Services,” on page 41
- ♦ Chapter 6, “Managing the Archive Server,” on page 49
- ♦ Chapter 7, “Configuring Jobs in iManager,” on page 53
- ♦ Chapter 8, “Managing Jobs,” on page 71
- ♦ Chapter 9, “Running Archive and Version Services in a Virtualized Environment,” on page 85
- ♦ Chapter 10, “Configuring Archive and Version Service for Novell Cluster Services,” on page 87
- ♦ Chapter 11, “Migrating Novell Archive and Version Services to OES 11,” on page 97
- ♦ Chapter 12, “Security Considerations for Archive and Version Services,” on page 99
- ♦ Appendix A, “XML Elements and Attributes for ArkConfig,” on page 101
- ♦ Appendix B, “Sample Configuration Files,” on page 117
- ♦ Appendix C, “Troubleshooting,” on page 119

Audience

This guide is intended for network administrators.

Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comments feature at the bottom of each page of the online documentation.

Documentation Updates

For the most recent version of the *OES 11: Novell Archive and Version Services 2.1 Administration Guide*, see the latest [Novell Open Enterprise Server 11 documentation \(http://www.novell.com/documentation/oes11/index.html\)](http://www.novell.com/documentation/oes11/index.html).

Additional Documentation

For documentation on accessing and restoring archived file versions, see the *OES 11: Novell Archive and Version Services 2.1 User Guide*.

For information about Novell Storage Services, see the [*OES 11: NSS File System Administration Guide for Linux*](#).

For information about Novell iManager, see the [Novell iManager 2.7.x Web site \(http://www.novell.com/documentation/imanager27/\)](http://www.novell.com/documentation/imanager27/).

1 Overview of Archive and Version Services

Novell Archive and Version Services 2.1 for Novell Open Enterprise Server 11 (OES 11) server provides a convenient and cost-effective way for individual users to instantly restore previous versions of modified, renamed, or deleted network files. It helps to minimize the user's unproductive time and frees you to focus on other critical IT issues. The user simply views a list of previous interval-based versions of source files, selects the file needed, then recovers it. The user can recover any of the available versions. If users work in a collaborative environment, a user can determine which version to choose based on who modified a document and when.

Archive and Version Services is available to archive user network files that are stored on Novell Storage Services volumes on OES 11 servers.

This section discusses the following:

- [Section 1.1, "Individual File Losses Impact Business," on page 9](#)
- [Section 1.2, "Benefits of Archive and Version Services," on page 10](#)
- [Section 1.3, "Key Concepts of Archive and Version Services," on page 11](#)
- [Section 1.4, "Scenarios for Using Archive and Version Services," on page 15](#)
- [Section 1.5, "What's Next," on page 17](#)

1.1 Individual File Losses Impact Business

Most enterprises implement some type of data backup and recovery to prevent major data losses. Backups occur periodically to prevent catastrophic losses of data. Often, the files that individuals lose have a life cycle shorter than the major backup cycles. Until now, these data losses have been an unfortunate cost of doing business.

Recovery of a single file is not usually a simple process. Only the administrator can access the backup media to retrieve and recover the file. The user must know exactly when the file existed so that the administrator can find the right version of the file. Even after the file is recovered, the user must update the file with changes made between the time it was backed up until the time it was modified, deleted, or lost.

Individual losses of key data impact business. However, most enterprises leave prevention and recovery to the best practices and personal habits of users. In a typical network environment, users employ different techniques to ensure that they do not lose critical files. For example, some users manually save multiple versions of a file under different names. Others save the same version of a file in different locations. Some do both.

Despite precautions, almost every user has accidentally modified, lost, or deleted a key file. When problems occur, the user is left with two choices:

- ♦ Wait for the administrator to recover the file from backup media, if the file was backed up at all
- ♦ Painstakingly rebuild the file from a backup version or from scratch

Either solution negatively impacts business:

- ♦ It's inefficient. The user cannot access backup files without administrator action.
- ♦ It's inconvenient. The user must waste time re-creating materials.
- ♦ It can affect the enterprise's ability to meet business commitments. Time lost can impact the user's ability to meet milestones, thereby impacting delivery to other processes down the line.

1.2 Benefits of Archive and Version Services

Novell Archive and Version Services for OES 11 provides benefits for the enterprise, IT administrators, and users.

Benefits for the Enterprise

Novell Archive and Version Services for OES11 offers two key benefits for the enterprise:

- ♦ It provides a lower cost of management for IT departments by allowing users to self-restore files from an archive of interval-based file versions.
- ♦ It provides a means to allow users to be more productive by allowing them to correct their own accidental deletions or file-modification mistakes.

Benefits for IT Administrators

For IT administrators, Novell Archive and Version Services for OES 11 solves the problem of individual file recovery. No longer does the Help Desk need to deal with users asking for a particular file to be restored. This frees IT organizations to focus on more important solutions for the users and the company as a whole.

Benefits for Users

With Novell Archive and Version Services, the user controls the file recovery; there is no need to completely rebuild a file or to involve the IT department. Users can retrieve file versions from anywhere, at any time, using a Web browser and an active network or Internet connection.

Novell Archive and Version Services for OES 11 offers many benefits for users:

- ♦ The versioning process is transparent to users until they need to retrieve a previous version of a file. Versioning does not affect how applications behave and requires no action on the part of the user.
- ♦ All security features and permissions of the source file are in effect for its file versions if the file versions are restored to an NSS volume on a OES 11 server. Using NetStorage, users can also download a file version as a new file, without its prior rights and metadata, to other types of storage media, such as to their local workstations.

- ♦ Versioning supports collaborative work environments where groups of users can create and modify shared files. The archive server allows a work group to properly select previous versions of files they are working on, based on who modified a file and the time stamp of the version.
- ♦ Novell Archive and Version Services for OES 11 uses Novell NetStorage to provide a Web-based interface to users for file version retrieval and restoration. Users can retrieve file versions from any workstation operating platform, including Microsoft Windows 2000/ XP, Windows Vista, Apple Macintosh, SUSE Linux Enterprise Desktop (SLED) 10 and later. All a user needs is a Web browser and an active network or Internet connection.
- ♦ Novell Archive and Version Services for OES 11 provides the NSS File Version Utility for Windows clients. Users can retrieve file versions and restore them from compatible Windows workstations.

1.3 Key Concepts of Archive and Version Services

Novell Archive and Version Services for OES 11 saves versions of user files at scheduled intervals, stores file versions in an archive database, and makes file versions available on demand to users. Novell Open Enterprise Server installs Novell Archive and Version Services by default as part of the OES 11 install. However, the service does not run until you configure and start it. For information, see [Chapter 5, “Setting Up Archive and Version Services,” on page 41.](#)

It is important to understand several key concepts and tools:

- ♦ [Section 1.3.1, “The Archive Server,” on page 11](#)
- ♦ [Section 1.3.2, “The Archive Database,” on page 12](#)
- ♦ [Section 1.3.3, “NSS Volumes,” on page 13](#)
- ♦ [Section 1.3.4, “Versioning Jobs,” on page 13](#)
- ♦ [Section 1.3.5, “Job Schedules,” on page 14](#)
- ♦ [Section 1.3.6, “File Versions,” on page 14](#)
- ♦ [Section 1.3.7, “Delete Policy,” on page 14](#)
- ♦ [Section 1.3.8, “Archive Versioning Plug-In for Novell iManager,” on page 14](#)
- ♦ [Section 1.3.9, “Versioning Policy,” on page 15](#)
- ♦ [Section 1.3.10, “NSS File Version Utility,” on page 15](#)
- ♦ [Section 1.3.11, “NetStorage Archive Function,” on page 15](#)

1.3.1 The Archive Server

The archive server runs Novell Archive and Version Services for OES 11, which includes the following services:

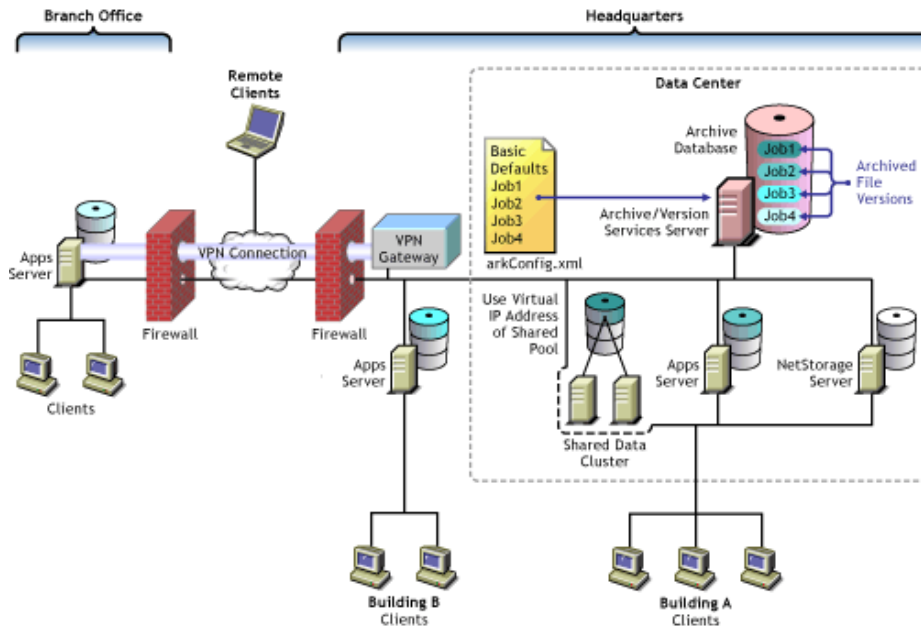
- ♦ Allowing users to search and restore file versions
- ♦ Versions of files based of configurable epoch intervals
- ♦ Non-intrusive backup of files
- ♦ Full access control, Workgroup based access
- ♦ Increased productivity

An archive server and the volumes resides in the same Novell eDirectory tree.

Novell Archive and Version Services supports several storage topologies, as shown in the following figure. Because file versions are transferred in decrypted format, the archive server should reside behind the corporate firewall. Any transfer of files during file versioning or restoration should occur over a secure connection such as a virtual private network (VPN). The volumes with files to be versioned can reside on local or remote servers and in single or clustered configurations.

For information, see [“Planning for Archive and Version Services” on page 21](#).

Figure 1-1 Example of Storage Topologies Supported by Novell Archive and Version Services



ArkManager

ArkManager is the software component that manages the file version process. This guide refers to Archive server as the server which runs the arkManager software.

For each archive server, you must configure the server’s basic properties, optionally configure default job settings, and configure properties for one or more individual jobs. Basic properties include details about the archive server and database. Default job properties specify the property settings to use instead of property settings specific to a particular job. An individual job specifies the property settings to use when archiving file versions for a specified volume that resides in the same eDirectory tree as the archive server.

The Archive Versioning plug-in for Novell Archive and Version Services allows you to configure server, defaults, and job settings and to manage jobs. For information, see [“Configuring Jobs in iManager” on page 53](#).

1.3.2 The Archive Database

Novell Archive and Version Services 2.1 uses a PostgreSQL database. PostgreSQL is an object-relational database management system (ORDBMS) based on POSTGRES package.

The Archive server stores the metadata information of the versions of the files in the PostgreSQL database. The database is known as the Archive database and resides on the Archive server.

1.3.3 NSS Volumes

- ♦ **Primary Volume:** The Primary volume contains the user files to be versioned.
- ♦ **Archive Volume:** The Archive volume contains the versions of the user files.

1.3.4 Versioning Jobs

A versioning job captures copies of eligible files on a specified source volume at specified intervals. Eligible files are those that exist in the source volume at the time the volume is versioned, meet the general versioning criteria, and pass any administrator-specified filtering criteria. You can define only one job for a given source volume.

Each job identifies the settings for the following properties. For details, see [Section 3.4, “Understanding Job Properties,”](#) on page 25.

Table 1-1 Overview of Job Properties

Property	Description
Name	The unique, administrator-specified job name that represents the relationship between the archive server and a given source volume. The job name persists for the life of the archive server and can represent only the specified source volume.
Server	The OES 11 server where the data to be versioned is located.
Mount Point	The OES 11 NSS volume where the data to be versioned is located. Each volume can be the target of only one versioning job.
Snapshot Pool	The OES 11 NSS pool where the snapshots of the source volume are temporarily stored while file versions are written to the archive database.
Free Space ID	ID of the free space object to be used for storing snapshot data. For example, /dev/hda.
Sectors	Specifies the number of sectors on free space to be used for storing snapshot data. The value needs to be an integer, for example, 409600.
Filter	Sets criteria to determine which files in the source volume are eligible for versioning.
Run Schedule	Specifies when to start the job and the frequency for running the job
Delete Policy	Determines the lifetime of file versions in the archive.

Archive Version Services is not supported on iFolder and Novell Distributed File Service.

1.3.5 Job Schedules

You must establish a schedule for each versioning job that meets users' requirements for file versioning, given limited storage and bandwidth resources. Versioning occurs for eligible files at scheduled intervals, called epochs. In Novell iManager, you can also manually pause versioning jobs and run jobs on demand, as needed.

A file's lifetime must span the end of an epoch to be versioned. Only files that exist when the versioning occurs are eligible to be versioned. If a user creates and deletes a file within the epoch, it cannot be versioned.

For more information, see [Section 3.4.3, "Run Schedule," on page 27](#).

1.3.6 File Versions

File versions are actual copies of files taken at scheduled intervals, as determined by the administrator. No matter how many changes users makes to files during an epoch, only those eligible files that exist at the end of the epoch are saved.

Novell Archive and Version Services 2.1 for OES 11 can use NSS pool snapshot technology to capture point-in-time copies of all files, even if the file is in use when the versioning process begins. If the snapshot option is not used, the versioning process captures only eligible files that are not deleted and not exclusively opened at the time.

User needs and limited storage and bandwidth resources are key considerations for setting the criteria to determine which files are eligible for versioning. Files can be filtered to include or exclude source files, according to their path, file extension, or filename patterns. If a user's files meet the filtering criteria, they are eligible for versioning. For information, see [Section 3.4.5, "Filter," on page 29](#).

Users do not have direct control over which of their files get versioned, when the versioning occurs, or the state of their files when the epoch ends and the copy is made. Users can access files natively with the NSS File Version Utility on a SLED 10/ SLED 10 SP1 or Windows 2000/ XP/ Vista desktop, or they can access their file versions at any time and from anywhere using the NetStorage Archive function. For information, see the [OES 11: Novell Archive and Version Services 2.1 User Guide](#).

1.3.7 Delete Policy

The Delete Policy determines when and which of a job's file versions to automatically delete from the archive database. Versioned files can have a limited lifetime in the archive. You optionally configure a job's Delete Policy to set the maximum keep time and the maximum number of versions to retain. The Delete Policy can allow indefinite retention of at least one most recent versioned file.


For more information, see [Section 3.4.4, "Delete Policy," on page 28](#).

1.3.8 Archive Versioning Plug-In for Novell iManager

After you configure your archive server and its versioning jobs, you can use the Archive Versioning plug-in for Novell iManager to manage those jobs. You can start and stop jobs, view a list of jobs, and view information about jobs, such as their current status, configuration details, and run schedules. You can also view the ArkManager log, which lists all normal, warning, and error messages for each job on the archive server.

For information, see [Section 7.1, "Accessing the Archive Versioning Plug-In in iManager," on page 53](#).

1.3.9 Versioning Policy

The Versioning Policy determines the criteria to filter the type of files to be versioned, duration and number of versions retained for a given volume. To access this information, the user right-clicks the versioning icon  in the system tray and selects *View Versioning Policy*.

1.3.10 NSS File Version Utility

The NSS File Version Utility provides convenient and direct access in a native Windows 2000/ XP/ Vista or SLED 10/ SLED 10 SP 1 environment to archived versions of user files. The utility integrates with your desktop to provide a *Restored Archived Versions* option, which allows users to view recent versions of their files and restore the desired file. Users select the desired version of the file, then click *Restore* to view the file locally or to restore the file version to a network storage location.

The NSS File Version Utility, is provided on the Welcome page. For information about using the utility, see the [OES 11: Novell Archive and Version Services 2.1 User Guide](#).

1.3.11 NetStorage Archive Function

File versions reside in the archive database on the archive server. Users can restore file versions from the archive database at any time from anywhere using the Archive function in Novell NetStorage. Using the NetStorage interface in the enterprise portal, a user views a list of available versions of a file. The user simply selects the previous version of the file, then clicks *Restore* to download the file version to a specified location where the user has the necessary permissions.

If a user restores the file version to a OES 11 NSS storage location, the archive server recovers the file version and all the rights and metadata about the file. If a user opts to download the file version elsewhere, the file is saved as a new file, without the prior rights and associations.

For information, see the [OES 11: Novell Archive and Version Services 2.1 User Guide](#).

1.4 Scenarios for Using Archive and Version Services

Novell Archive and Version Services 2.1 for OES 11 is a valuable asset in today's diversified workplace. This section discusses the following usage scenarios:

- [Section 1.4.1, "Cross-Platform Work Environments," on page 15](#)
- [Section 1.4.2, "Group Collaboration," on page 16](#)
- [Section 1.4.3, "File and Directory Name Changes," on page 16](#)
- [Section 1.4.4, "Selective File Versioning," on page 16](#)

1.4.1 Cross-Platform Work Environments

Novell Archive and Version Services for OES 11 provides a Web-based interface to the archive database in Novell NetStorage with the Archive function. Users can retrieve file versions from any workstation operating platform, including NetWare, Microsoft Windows 2000/XP / Vista, Apple Macintosh, SLED 10, and SLED 10 SP1. All a user needs is a Web browser and an active network or Internet connection.

1.4.2 Group Collaboration

In a shared work group, a team works collectively to share information, create information, and process information. Files are regularly shuffled back and forth between users, and they are usually worked on by more than one person. Sections of a presentation are created by different people and either merged or are simply edited into an existing file that is passed around.

Research has shown that people tend to solve the problems associated with lack of versioning in PC network systems by attempting to do an ad hoc versioning system. Unfortunately, everyone does it a bit differently. Some put a version or date in the name of the file, such as `MarketAnalysisv3.ppt`. Others put dates in the filename, such as `MarketAnalysis2002Oct03.ppt`. Still others use file folders with versions or dates in the names of the folders. Most are not consistent with their techniques and many do not even try. It is especially troublesome when the files are shared, because not every personal scheme is alike. Even with these various methods, mistakes happen.

Novell Archive and Version Services for OES 11 supports collaborative work environments. The archive interface allows a user to view previous versions and see instantly who was the modifier of each of the versions without opening file versions to attempt to ascertain who modified it.

For example, Tom and Alice worked together to prepare a presentation. Two weeks ago, Alice deleted some edits that Tom made to the file. Now, the team needs those edits back. Alice cannot recall when she deleted the edits. By going to the Web-based archive access, Alice can view previous file versions. The modifier of the file is listed next to each file version. Alice easily identifies the file version from about two weeks ago that shows up with Tom as the modifier of the file. She can view the file version or restore the file, as needed, to recover the lost modifications.

1.4.3 File and Directory Name Changes

File and directory names are likely to change during their lifetimes. Novell Archive and Version Services for OES 11 supports file and directory renaming. It tracks changes to the filename, such as when a user renames a file at some point in the process of its creation and modification. It also tracks changes made to the file's subdirectory (or path) if it is changed.

1.4.4 Selective File Versioning

It is not desirable or practical to version every file in a volume. Novell Archive and Version Services allows administrators to determine which files on their servers get versioned and the versioning intervals on a per-volume basis.

For example, consider a source volume that comprises multiple directories: Users, Shared, and several directories for applications. Although the files in the Users and Shared directories change frequently, the files in the applications directories are fairly stable. Novell Archive and Version Services allows the administrator to exclude files in applications directories from versioning.

The administrator can selectively control the versioning frequency for each volume. For example, consider a Users volume with files that change intermittently throughout the day and a Shared volume with files that change at the end of each scheduled work shift. Novell Archive and Version Services allows the administrator to schedule 30-minute epochs for the Users volume and schedule the start time and subsequent epochs to coincide with shift changes for the Shared volume.

The administrator can control the types of files to be versioned. For example, consider a Productivity volume that contains both Web development applications and user files. With Novell Archive and Version services, the administrator can specify the extensions of files like `.doc` or `.html` to be versioned in a Productivity volume.

1.5 What's Next

Use the following table to determine where to find information:

Table 1-2 *Possible Tasks*

To Perform This Task	Refer To
Planning your Novell Archive and Version Services 2.1 for OES 11 implementation	Planning for Archive and Version Services (page 21)
Assessing your implementation plan against the prerequisites and guidelines	Prerequisites and Guidelines (page 33)
Setting up a new archive server	"Setting Up Archive and Version Services" on page 41
Configuring the versioning jobs	Configuring Jobs in iManager (page 53)
Managing the archive server	Managing the Archive Server (page 49)
Managing the versioning jobs	Managing Jobs (page 71)
Managing security for archive services	Security Considerations for Archive and Version Services (page 99)

2 What's New or Changed in Archive and Version Services

The Archive and Version Services 2.1 service has been modified to run on OES 11. There are no other changes in the OES 11 release of Archive and Version Services 2.1.

3 Planning for Archive and Version Services

This section discusses how to plan and design Novell Archive and Version Services 2.1 for OES 11 to meet your business needs.

- ♦ [Section 3.1, “Assessing Your Versioning Needs,” on page 21](#)
- ♦ [Section 3.2, “Designing the Archive and Version Services Topology,” on page 23](#)
- ♦ [Section 3.3, “Understanding Archive Server Properties,” on page 24](#)
- ♦ [Section 3.4, “Understanding Job Properties,” on page 25](#)
- ♦ [Section 3.5, “What’s Next,” on page 32](#)

3.1 Assessing Your Versioning Needs

Before implementing Novell Archive and Version Services 2.1 in your network environment, collect information about your system to assess your versioning needs. Ask and answer the following questions:

- ☐ **Users:** Which users would benefit from having an archived database of multiple historical versions of their network files? For example, in a university environment, you might provide versioning support for faculty and staff, but not for campus lab environments.

Examine business and operational activities to identify and prioritize users’ versioning needs. Use this information to plan a strategic implementation of versioning throughout your network.

- ☐ **Volumes and Directories:** What volumes do these users use? How are the user directories organized on each volume?

If you are versioning only selected users’ data on a volume, you can exclude all data, and then include each user’s data by path.

If you archive files from an encrypted volume, the destination path for Archive Manager should also be on an encrypted volume. If the destination path is a nonencrypted volume, the versioned data is stored in a nonencrypted state.

For any source volume, you can define only one versioning job. If you attempt to define multiple jobs for a volume on the same or different server, Archive and Version Services does not run as designed and data integrity in the archive database is compromised.

- ♦ The archive server can be the same or different server as the source server.
- ♦ A single archive server can archive many volumes with only one job defined per volume.

- ♦ A given volume cannot have multiple jobs defined for it, even if the jobs run on different archive servers.
- ♦ Files and directories that reside under Distributed File Service (DFS) junctions are not versioned, because Archive and Version Services does not support junctions. If you need to archive data that resides on the target volume of that junction, set up jobs on the target volume and server pointed by that junction.

- ❑ **File Extensions:** What types of data within these volumes are candidates for versioning? For example, productivity files such as documents, spreadsheets, presentations, graphics, and other file types typical in your industry.

Many types of data do not need versioning and can quickly consume valuable space in your archive database. Avoid versioning system files, log files, and databases. In addition, do not version temporary files that have no value to users, such as temporary Internet files and temporary application files. Identify the extensions of temporary files to be excluded from the job.

Archive and Version Services fails to create versions for iFolder or any other application that creates backup files.

- ❑ **Versioning Frequency:** How often does the identified data change on each volume? You should set the frequency to best satisfy the needs of the users. Look for patterns in frequency, such as the following:
 - ♦ Does frequency vary by types of data?
 - ♦ Does frequency vary by individual users?
 - ♦ Does frequency vary by categories of user, such as power users, support groups, or manufacturing?
 - ♦ Does frequency vary by physical location? For example, headquarters, branch offices, geographical regions, or departments.
 - ♦ Does frequency vary by time? For example, by time of day, day of the week, week in the month, month in the year, fiscal quarter, season (winter, spring, summer, or fall), or special events.
- ❑ **Storage:** What are the minimum initial storage capacity and scalability needs for your archive database? Determine this information by assessing your current storage requirements. For example:
 - ♦ How much storage space is consumed by each data type? For example, estimate the number of files and their average size.
 - ♦ What is the expected growth of data? For example, the additional number of megabytes or gigabytes of storage needed by users by month, quarter, or year.
 - ♦ What is the rate of growth in storage capacity by data type? For example, evaluate the percentage increase over time of a particular file type, such as documents, spreadsheets, presentations, graphics, or other file types typical of your industry.
- ❑ **Archive Database:** The potential size of the archive database depends on the combination of the following:
 - ♦ How many versioning jobs are defined for the server?
 - ♦ How often does each job run?
 - ♦ How many files are versioned on the source volume?

- ♦ How often are the files modified and how does that correspond to the scheduled interval for the versioning epoch?
 - ♦ What is the delete policy for the job?
- ❑ **Topology:** Where in the enterprise network are the volumes stored? For example, a data center, a branch office, or work islands throughout the campus.
- ❑ **Delete Policy:** What type of retention policy for file versions do you need for each volume? Establishing a delete policy helps you control the storage capacity required by the archive database. For example:
- ♦ How long do you need to retain versions?
 - ♦ What is the maximum number of versions to keep, given the number of source files versioned and their sizes?
 - ♦ Do you need to keep at least one file version after its source file has been deleted?
 - ♦ Do you need a complete copy of all the data to begin the archive database, or do you want the collection of file versions to grow only as data changes?

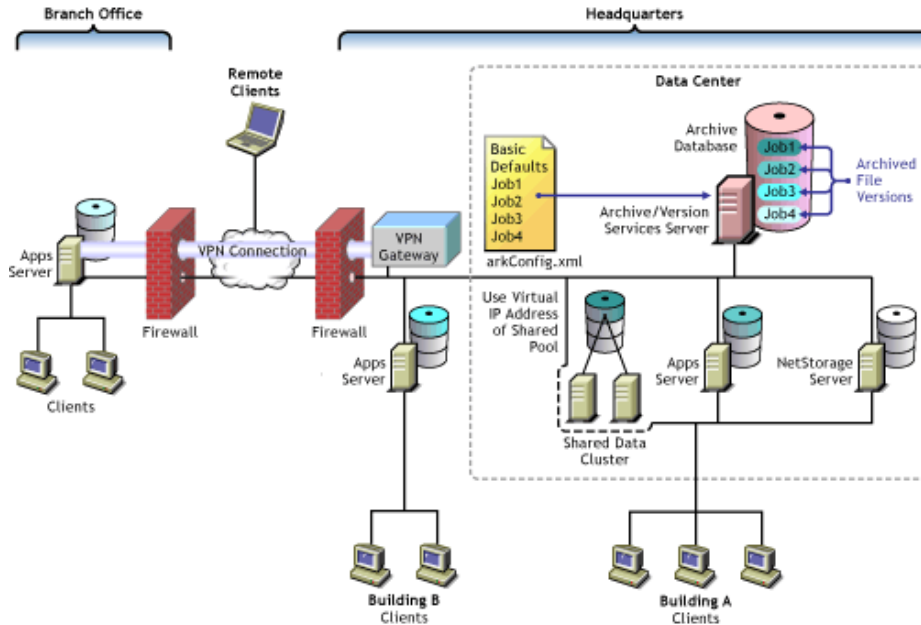
3.2 Designing the Archive and Version Services Topology

Novell Archive and Version Services 2.1 supports several storage topologies, as shown in the following figure. In a typical enterprise, the archive server resides in the data center to provide convenient access to services for all users. The source volumes with files to be versioned can reside on local or remote servers in single or clustered configurations.

The number of archive servers you need depends on the number and size of file versions being archived and the distribution of users who need to access the versions. For example, if you have numerous volumes with many files changing at a high frequency, you might implement multiple archive servers. You should also consider the feasibility of backup for a server based on the anticipated size of the archive database and data.

For a geographically distributed enterprise, an archive server might reside in each regional data center. Key factors to consider are the speed, capacity, security, and cost of communications links between regions. Because the archive server transfers files non-encrypted across network connections, communications environments and connections must be secure.

Figure 3-1 Example of Storage Topologies Supported by Novell Archive and Version Services



3.3 Understanding Archive Server Properties

The archive server properties define the basic information that applies to all jobs controlled by an archive server. They specify the authentication information for Novell eDirectory and the archive database, whether to display the log using iManager at *Archive Versioning > Archive Jobs > Log* tab and the storage location where the archive data (file versions) reside.

Configure the following server properties using Archive Versioning plug-in:

Table 3-1 Description of Properties for the Archive Server

Archive Server Properties	Description
Volume Path	Specifies the location of archive data in the OES 11. For example, Ark : / archive.
Display server Archive Log entries	<p>By default, the archive server records the error, warning, and normal messages for all of its jobs in the Archive Log available at /var/opt/novell/arkManager/logs file.</p> <p>In iManager if the <i>Display Log</i> option is selected, the log messages are displayed at <i>Archive Versioning > Archive Jobs > Log</i>.</p>

Archive Server Properties	Description
eDirectory	<p>Specifies authentication information about the Archive and Version Services administrator user.</p> <ul style="list-style-type: none"> ♦ User Name: The eDirectory Common Name of the administrator user who has the appropriate rights to the original data location and to the archive data location. For example, <code>admin.servercontext</code> The archive server administrator user must have the file system Supervisor right to the archive server and to all servers being accessed by the selected archive server. The user must be in the same eDirectory tree as the archive server and the source servers.
Database	<p>Specifies authentication information about the PostgreSQL database for the archive server.</p> <ul style="list-style-type: none"> ♦ User Name: The administrator user of PostgreSQL. For example, <code>postgresadmin</code>. ♦ Port: The port used by the ArkManager instance of PostgreSQL on the archive server. By default, Port 5432 is used. If Port 5432 is used by other services, you can configure other ports.

3.4 Understanding Job Properties

Job properties are the set of parameters used for individual job settings and default job settings. Individual job settings apply only to a single job. Default job settings can apply, in whole or in part, to any individual job defined for an archive server. Parameters used only for individual jobs include Name, Volume, and Stopped. Within a job definition, you can specify values for a given job property, or if you do not specify the value, its default value can be applied. In some cases, you can choose to specify no value, such as when a function is disabled.

Changing the property's value in an individual job's settings causes a different outcome than changing property's value for the default job settings. Modifying values in an individual job affects only the behavior of the individual job. Modifying values for the default job settings affects every job that uses the default values. All changes take effect the next time ArkManager runs.

Use the following properties to define jobs for your archive server:

- ♦ [Section 3.4.1, "Job Information," on page 25](#)
- ♦ [Section 3.4.2, "Source Server Information," on page 26](#)
- ♦ [Section 3.4.3, "Run Schedule," on page 27](#)
- ♦ [Section 3.4.4, "Delete Policy," on page 28](#)
- ♦ [Section 3.4.5, "Filter," on page 29](#)

3.4.1 Job Information

The Job Information identifies control information for jobs on an archive server. Each job has a unique name that represents a unique relationship between an archive server and a source volume.

IMPORTANT: You can define only one job per source volume.

The following table describes the job information properties:

Table 3-2 *Description of Job Properties for Job Information*

Job Property	Description
Name	<p>The administrator-specified unique job name. For example, svr1_users or svr2_finance. A job name can be up to 64 ASCII characters.</p> <p>The job's name persists for the life of the archive server and represents the relationship between the archive server and a specific source volume. You cannot reuse the name for any other archive server and volume relationship.</p> <p>Specify the Name property only for individual jobs, not as a default value.</p>
Stopped	<p>Specify this property to define the job but leave it in a Stopped state until you manually activate the job. If the <i>Stopped</i> property is not used, the job starts, according to the Run Schedule settings, the next time ArkManager runs.</p> <p>Specify the <i>Stopped</i> property only for individual jobs, not as a default value.</p>
Full Copy or No Full Copy	<p>If Full Copy is selected, the archive server copies all files specified by the job from the source volume to the archive volume on a one-time, special version occurrence. Use this option to save at least one version of every file eligible for versioning. This copies the files the first time job is run.</p> <p>If No Full Copy is selected, the job begins at the next scheduled start time. Use this option to save versions of files eligible for versioning only as the files change during epochs.</p>

3.4.2 Source Server Information

The Source Server information identifies the source server, the NSS volume that contains the information to be versioned, and the destination pool where temporary snapshot pools are stored during the versioning process.

Table 3-3 *Description of Job Properties for Source Server Information*

Job Property	Description
Server	The host name of the OES 11 source server where the data to be versioned is located. For example, svr2. This server is in the same eDirectory tree as the archive server.
Mount Point	The OES 11 NSS volume where the data to be versioned is located. Each volume can be the target of only one versioning job. For example, primevol2.
Snapshot Pool	<p>The OES 11 NSS pool where the snapshots of the source volume are temporarily stored while file versions are written to the archive database.</p> <p>Specify the pool name of the destination pool for the snapshots. This pool is where the primary volume resides. The snapshots are maintained temporarily at the end of an epoch until point-in-time file versions can be saved to the archive database. For example, primepool2.</p> <p>If no snapshot pool name is specified or if the snapshot cannot be created for any reason, the versioning process copies files directly from the source volume. In this case, eligible files that are exclusively opened at the time cannot be versioned, and eligible files that are deleted before the copy occurs cannot be archived.</p>
Free Space ID	ID of the free space object to be used for storing snapshot data. For example, hda_nwfreespace1 or hda.
Sectors	Specifies the number of sectors on free space to be used for storing snapshot data. The value needs to be an integer, for example, 419430400.

3.4.3 Run Schedule

Run Schedule specifies when to start the job upon activation, and the frequency for running the job. To set the frequency, you must specify one of three scheduling options: Use Defaults, Scheduled Interval, or Scheduled Start Time.

Table 3-4 Description of Job Properties for the Run Schedule

Job Property	Description
Scheduled Interval	<p>If it is used, the Scheduled Interval specifies the elapsed time between the beginning of versioning processes for a job. The units can be seconds, minutes, hours, or days. For example, 45 seconds, 1 minute, 15 minutes, 2 hours, or 12 hours.</p> <p>If the versioning process exceeds the time specified as the interval, the overlapping scheduled job is skipped. No file versions are saved for skipped job runs. After the version process completes, the job runs at its next scheduled interval. If you observe the job skipping some versioning intervals, you can increase the interval between versions or reduce the amount of data to be versioned by setting Filter properties.</p>
Scheduled Start Time	<p>If it is used, the Scheduled Start Time specifies the start time when the job's version process begins and one or more days of the week to run the job.</p> <ul style="list-style-type: none">♦ Start: Specify the start time in hours and minutes.♦ Daily: Specify the days of the week to run the job. For example, specify All days for a daily job run, or specify one or more days of the week you want the job to run. Choices include Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, Sunday, and All.

3.4.4 Delete Policy

The Delete Policy determines the retention of file versions by age or by number of versions.

Table 3-5 Description of Job Properties for the Delete Policy

Job Property	Description
Delete Interval	<p>Specifies the schedule for deleting file versions, if they are eligible for deletion. The interval represents the amount of time to wait from the time a Delete process ends until another Delete process begins. If a value is not specified, 24 hours is the default interval.</p> <p>The time involved in deleting file versions varies with the amount of data stored in the archive server. The Delete Schedule operates separately from the Version Schedule.</p> <p>For example, suppose you set the Delete Schedule to 1 hour. When you activate the job, the Delete Schedule begins an interval timer. After 1 hour elapses, the Delete process runs. The timer is inactive while the process runs. When the delete process ends, the interval timer begins again. The process repeats until the job is deactivated.</p>

Job Property	Description
Maximum Keep	<p>Specifies the maximum number of versions of each file to keep in the archive and how long to keep file versions. At least one of the values must be non-zero. If you set both the Maximum Keep Versions and Maximum Keep Time to zero values, the Delete Policy function does not run.</p> <ul style="list-style-type: none"> ♦ Time: The maximum time that a file version is maintained in the archive. Specify the time in whole numbers for seconds, minutes, hours, or days. ♦ Versions: Specify the maximum number of versions of each file to keep in the archive. When the number of versions exceeds this integer value, the oldest version is deleted. ♦ Keep latest version of current file: By default, Keep latest version of current file option is enabled. Enabling this option makes the archive server keep an existing file version as long as its source file is current on the source volume, beyond the Maximum Keep Time. After the user deletes the current source file, the deletion is noted at the next scheduled epoch. If the file version's age is within the Maximum Keep Time, the archive database retains a copy of the file version until its Maximum Keep Time elapses. When the file version's age exceeds the Maximum Keep Time, the archive deletes the file version at the next scheduled delete time. <p>If Keep Current Copy is disabled, the archive deletes the file version when the Maximum Keep Time elapses.</p>

3.4.5 Filter

The Filter information determines what data in the source volume gets versioned. You can combine the filters for the individual job with filters for the default job settings.

You can filter files in the source volume that you do not want to version by using a series of Include and Exclude elements under the Filter section using Archive Versioning plug-in in iManager. The order of the Include and Exclude elements determines what data is eligible for versioning. Make sure the order is adequate to achieve the desired filtering outcome.

Filtering is optional, but you should avoid versioning volumes that contain system software. Exclude system files and file types that change constantly, such as log files and databases. You can also exclude nonessential file types such as MP3 and temporary files such as Internet files. Identify the file types your applications use as intermediate saves for open files, such as the TMP files for Microsoft Word, and set up filters to exclude that file extension from versioning.

Extremely large files, such as database files and ISO image files, take a long time to be copied into the archive, which can potentially block other requests to access the database. You cannot filter files by file size, but you can modify database settings or distribute data to lessen the impact of versioning large files. Another option is to separate larger files into one or more separate volumes, and then create a job for each source volume. Schedule the jobs to run in off-peak hours.

Table 3-6 *Description of Properties for the Filter*

Job Property	Description
Path	Specifies the relative path of directories in the specified source volume that you want to include or exclude in the versioning process. For example: <code><path>\log</path></code>
Extension	Specifies the extension of files in the specified source volume that you want to include or exclude in the versioning process. Use the preceding dot followed by the characters of the file extension. For example: <code><extension>.mp3</extension></code>
Pattern	Specifies the regular expression pattern to match for files that you want to include or exclude in the versioning process. For example, to specify a pattern for files that start with the letter a: <code><pattern>.*\a.*</pattern></code>
Wildcard	Specifies the wildcard pattern to match for files that you want to include or exclude in the versioning process. For example: <code><wildcard>a*tmp</wildcard></code>

Include and Exclude Elements

You can filter out files in the source volume that you do not want to version by using a series of Include and Exclude elements. The child elements within each Include or Exclude element can contain multiple Path, Extension, and Pattern elements, in whatever order is needed to determine what data is eligible for versioning.

By default, all data in the volume is included. The first step in filtering is to exclude everything. For example:

```
<exclude>
  <path>\</path>
</exclude>
```

Next, add back in the paths, file types, and patterns for files you want to version. The nested include or exclude elements always override previous include or exclude elements.

Pattern Elements

The regular-expression parser used for the `<pattern>` tag does not support the following regular-expression constructs in PERL 5:

- The conditional constructs `(?{X})` and `(?(condition)X|Y)`
- The embedded code constructs `(?{code})` and `(??{code})`
- The embedded comment syntax `(?#comment)`
- The preprocessing operations `\l`, `\u`, `\L`, and `\U`

The regular-expression parser used for the <pattern> tag supports the following regular-expression constructs, which PERL 5 does not:

- ♦ Possessive quantifiers, which match as much as they can and do not back off, even when doing so would allow the overall match to succeed
- ♦ Character-class union and intersection

Character classes can appear within other character classes, and can be composed by the union operator (implicit) and the intersection operator (&&). The union operator denotes a class that contains every character that is in at least one of its operand classes. The intersection operator denotes a class that contains every character that is in both of its operand classes.

The precedence of character-class operators is as follows, from highest (1) to lowest (5):

- 1) Literal escape \x
- 2) Grouping [. . .]
- 3) Range a-z
- 4) Union [a-e][i-u]
- 5) Intersection [a-z&&[aeiou]]

Other notable differences from PERL-based regular expressions are shown in the following table.

Table 3-7 Comparison of Supported Regular Expressions and PERL-Based Regular Expressions

Supported Regular Expressions	PERL-Based Regular Expressions
Octal escapes must always begin with a zero. \1 through \9 are always interpreted as back references, and a larger number is accepted as a back reference if at least that many sub-expressions exist at that point in the regular expression; otherwise, the parser drops digits until the number is smaller or equal to the existing number of groups or it is one digit.	\1 through \9 are always interpreted as back references; a backslash-escaped number greater than 9 is treated as a back reference if at least that many sub-expressions exist; otherwise, it is interpreted, if possible, as an octal escape.
It is implicit that repeated invocations of the find method resume where the last match left off, unless the matcher is reset.	PERL uses the g flag to request a match that resumes where the last match left off.
Embedded flags always take effect at the point where they appear, whether they are at the top level or within a group; in the latter case, flags are restored at the end of the group just as in PERL.	In PERL, embedded flags at the top level of an expression affect the whole expression.
The defined class accepts dangling brackets but is strict about dangling metacharacters like +, ?, and * throws a PatternSyntaxException if it encounters them.	PERL is forgiving about malformed matching constructs, as in the expression *a, as well as dangling brackets, as in the expression abc], and treats them as literals.

For more information, consult a programming textbook or search the Internet for a reference that discusses the behavior of regular expression constructs.

Wildcard Elements

A wildcard functions like a wildcard in directory searches. Replace characters with an asterisk (*) to search for files that match. For example, to include or exclude files that start with d of type .sxi:

<wildcard>d*sxi</wildcard>

3.5 What's Next

Before you deploy Novell Archive and Version Services 2.1, make sure your deployment plan satisfies the following:

- ♦ [Chapter 4, “Prerequisites and Guidelines,” on page 33](#)
- ♦ [Chapter 12, “Security Considerations for Archive and Version Services,” on page 99](#)

4 Prerequisites and Guidelines

This section discusses prerequisites and guidelines for designing your Novell Archive and Version Services 2.1:

- ♦ [Section 4.1, “Network Architecture Prerequisites and Guidelines,” on page 33](#)
- ♦ [Section 4.2, “Server Prerequisites and Guidelines,” on page 33](#)
- ♦ [Section 4.3, “Storage Media Prerequisites and Guidelines,” on page 34](#)
- ♦ [Section 4.4, “Fault Tolerance Guidelines,” on page 35](#)
- ♦ [Section 4.5, “Job Guidelines,” on page 35](#)
- ♦ [Section 4.6, “Schedule Guidelines,” on page 36](#)
- ♦ [Section 4.7, “Prerequisites for Using the NSS File Version Utility,” on page 36](#)
- ♦ [Section 4.8, “Prerequisites for Using the NetStorage Archive Function,” on page 37](#)
- ♦ [Section 4.9, “Prerequisites and Guidelines for Retrieving File Versions,” on page 37](#)
- ♦ [Section 4.10, “Guidelines for Availability of File Versions,” on page 38](#)
- ♦ [Section 4.11, “What’s Next,” on page 39](#)

4.1 Network Architecture Prerequisites and Guidelines

Make sure your network meets the following prerequisites:

- ♦ Because the archive server moves data in decrypted format, the archive server should be located behind the corporate firewall. If you are versioning data on remote servers, use a virtual private network (VPN) connection between the two.
- ♦ Determine how much bandwidth is required to support versioning traffic, given the following criteria:
 - ♦ How much data is transferred as file versions during each epoch
 - ♦ When the scheduled epochs occur
 - ♦ Where in the network the source volumes are located (the path between the archive server and the source server)
 - ♦ Peak traffic flow based on seasonal variations in productivity

4.2 Server Prerequisites and Guidelines

Make sure your archive server and source servers meet the following prerequisites and guidelines:

- ♦ The Archive server can be the same or different server as the source server.
- ♦ The Archive server and the Client workstation system clocks must be synchronized.

- ♦ One archive server can have only one job defined for any one volume.
- ♦ A single archive server can archive many volumes with only one job defined per volume.
- ♦ A given volume can have only one job defined for it.
- ♦ The archive server and the source server can be the same machine or different machines.
- ♦ The archive server and primary server both should be in the same eDirectory tree.
- ♦ The archive data volume must reside on the archive server.
- ♦ Make sure your server meets the system and software requirements for OES 11. For information, see the [OES 11: Installation Guide](#).
- ♦ Determine how much storage space you need for the archive database, given the following criteria:
 - ♦ How much data needs to be stored immediately by the versioning job
 - ♦ How much data there is to version (the physical storage requirements by versioning job)
 - ♦ How frequently the data must be versioned (the length of time in an epoch by versioning job)
 - ♦ How long you need to store the versions (maximum keep time and maximum versions to keep by versioning job)

4.3 Storage Media Prerequisites and Guidelines

Make sure your archive volume and source volumes meet the following prerequisites and guidelines:

- ♦ The source and archive data volumes must be Novell Storage Services (NSS) data volumes.

You can create the NSS volumes on an OES 11 server. If the NSS volume is on an OES Linux server, you must move it to a cross-compatible OES 11 server before it can participate in your Archive and Version Services solution.
- ♦ Create an NSS storage pool and volume to use exclusively for your archive database and archive data that meets the following guidelines:
 - ♦ For each archive server, use the information from your planning session to design the physical storage media for its archive database.
 - ♦ Make sure your implementation allocates enough space to meet your immediate archive needs and is scalable to accommodate future growth.
 - ♦ The NSS volume you use for your archive database cannot contain data that you plan to archive. The archive volume and the source volume cannot be the same volume.
 - ♦ The archive volume must be a local volume.
 - ♦ If you plan to version files from an encrypted source volume, the archive volume should also be an encrypted volume so that data continues to be protected in the archive database. For information about Encrypted Volume Support, see Managing Encrypted NSS Volumes in the [OES 11: NSS File System Administration Guide for Linux](#).
- ♦ We recommend that the archive volume and the source volume should reside in separate storage pools, but this is not a limitation.

- ♦ We recommend you to use different storage devices for the pool for the archive volume and the pool for the source volumes, but this is not a limitation.
- ♦ For easy backup, we recommend you to store the archive database and archive data in a directory in the archive volume. The archive data is always stored in subdirectories at the base of the archive path.

4.4 Fault Tolerance Guidelines

If your data is critical, you can design fault tolerant and high availability solutions for Novell Archive and Version Services, including multiple connection channels, software RAID devices, and cluster solutions. These solutions are optional.

4.4.1 Multiple Connection Channels for Storage Devices

Multiple connection channels can help ensure fault-tolerant connectivity between the archive server and the devices containing the archive volume.

For information, see the [OES 11: NSS File System Administration Guide for Linux](#).

4.4.2 Software RAID Devices

Software RAIDs can improve read/write performance and ensure data protection. NSS file systems support software RAID 0 (striping), RAID 1 (mirroring), and RAID 5 (striping with parity). All software RAID devices can improve file access performance, but only RAID 1 and RAID 5 also provide data protection to your storage media solution. To add fault tolerance your archive volume, you can set up the device you plan to use for your NSS pool as a RAID-1 device, a RAID-5 device, or a RAID-10 device.

For information, see Managing Software RAID Devices in the [OES 11: NSS File System Administration Guide for Linux](#).

4.4.3 Server Clusters

A Novell Cluster Services solution improves service availability. If a server goes down, the Cluster Services software handles the failover to the next available server in the cluster. If the server running Archive and Version Services goes down, Cluster Services ensures that the archive services and database remains available to versioning processes and to users who need to retrieve file versions.

You need to install Novell Cluster Services (NCS) before configuring Archive and Version Services to work in a cluster environment. NCS provides an Archive and Version Services resource template, which facilitates configuring the service in an cluster environment. For more information, see [Chapter 10, “Configuring Archive and Version Service for Novell Cluster Services,” on page 87](#)

4.5 Job Guidelines

One archive server can only have one job defined for any one volume. Multiple archive servers cannot define jobs on the same volume. However, a single archive server can run jobs on multiple source volumes.

WARNING: Define only one job per volume. If you attempt to define multiple jobs for a volume, ArkManager does not run as designed and data integrity in the archive database is compromised.

Jobs should focus on versioning files from productivity applications. Filtering is optional, but you should avoid versioning volumes that contain system software. Exclude system files and file types that change constantly, such as log files and databases. You might also want to exclude nonessential file types such as MP3 and temporary files such as Internet files. Identify the file types your applications use as intermediate saves for open files, such as the TMP files for Microsoft Word, and set up filters to exclude that file extension from versioning.

Extremely large files, such as database files and ISO image files, take a long time to be copied into the archive, which can potentially block other requests to access the database. You cannot filter files by file size, but you can modify database settings or distribute data to lessen the impact of versioning large files. Another option is to separate larger files into one or more separate volumes, and then create a job for each source volume. Schedule the jobs to run in off-peak hours.

A job name must be unique to the archive server. To see job names that are currently in use, see [Section 8.1, “Viewing a Jobs Report,” on page 71](#). However, after you delete a job, it no longer appears in this list.

4.6 Schedule Guidelines

You should allow enough time for one version job to be captured before beginning another epoch for that job. If a versioning process cannot finish before the next process is scheduled to begin, the archive server continues to version the job and skips other scheduled processes until the job finishes.

For example, suppose that you configure a versioning job with a one-minute epoch for a source volume that contains an amount of data that takes 5.2 minutes to archive. The version process begins at time 0. At each 1-minute interval thereafter, the next scheduled versioning checks to see if the job is already running. If it is, the job does not begin. In this example, the scheduled versioning processes do not occur at 1, 2, 3, 4, and 5 minutes. The next versioning process begins at 6 minutes. Although the user expects to have file versions available at 1-minute intervals, the file versions are at 6-minute intervals.

To avoid such problems, consider the following configuration alternatives:

- ♦ Set epochs realistically, based on the amount of data to be versioned from the source volume and the bandwidth of the connection between the source server and the archive server.
- ♦ Filter the data on the source volume to exclude unnecessary data from being versioned.
- ♦ Divide data on the source volume to create multiple volumes. Then you can configure multiple jobs to run concurrently on the same or different archive server.

4.7 Prerequisites for Using the NSS File Version Utility

Users can search for file versions and restore them using the NSS File Version Utility. The utility integrates with the Windows XP/ Vista or SLED 10 or SLED 10 SP 1 desktop to allow users to natively restore previous versions of their current, renamed, or deleted files that are stored in the archive database.

- ☐ NSS File Version Utility for Linux

Type	Description
Operating System	SLED 10 or later
Dependencies	<ul style="list-style-type: none"> ♦ Novell Client 1.2 SP2 ♦ KDE or GNOME base libraries ♦ Mono 1.2.2
File Manager	KDE or Nautilus

- ☐ NSS File Version Utility for Windows

Type	Description
Operating System	<ul style="list-style-type: none"> ♦ Windows 2000 Professional SP4 ♦ Windows XP SP2 ♦ Windows Vista
Dependencies	<ul style="list-style-type: none"> ♦ Novell Client 4.9.1 SP 2 or later ♦ Microsoft .NET 2.0 framework for Windows XP

- ☐ Set up NCP support for the archive server.

For information about installing and using the NSS File Version Utility, see the [OES 11: NSS File System Administration Guide for Linux](#).

4.8 Prerequisites for Using the NetStorage Archive Function

Novell NetStorage provides Web-based interface for retrieving files from the Archive server. Users can search for file versions and restore them with Novell NetStorage from anywhere, at any time, using a compatible Web browser and a network or Internet connection.

- ☐ Install a [OES 11: NetStorage Administration Guide for Linux](#).
- ☐ Configure NetStorage services to access Primary volume on the Primary server.

For information about installing and configuring Novell NetStorage, see the [OES 11: NetStorage Administration Guide for Linux](#).

For information about using NetStorage and the Archive function, see the [OES 11: NetStorage Administration Guide for Linux](#).

4.9 Prerequisites and Guidelines for Retrieving File Versions

Previous versions of files can exist in the archive database only after the files meet the following prerequisites:

- ♦ You install and configure Novell Archive and Version Services 2.1 on your OES 11 server.
- ♦ You configure at least one job to version OES 11 NSS volumes where user data is stored.

- ♦ You set criteria to determine which files are eligible for versioning. Files can be included or excluded according to their path, file extension, or filename patterns. If files on a specified volume meet the resultant criteria, they are eligible for versioning.
- ♦ Versioning occurs for eligible files at scheduled intervals, called epochs. You can also use iManager to pause jobs and run jobs on demand, as needed.
- ♦ A file's lifetime must span the end of at least one epoch to be versioned. Only files that exist when the versioning occurs are versioned.

If a user creates a file after an epoch begins, and deletes the file before the end of the epoch, the file cannot be versioned. If a user creates a file that spans the end of a scheduled epoch, but deletes the file before the scheduled job actually copies the file to the archive, the file is not versioned.

- ♦ It does not matter how much or how often a user changes an eligible file during an epoch. The versioning process captures the file in whatever state it is in at the end of the epoch. A user does not have direct control over which files are versioned, when the versioning occurs, or what the state of any file is when the epoch ends and the file versions are copied to the archive database.
- ♦ Versioned files typically have a limited lifetime in the archive. You configure a Delete Policy that sets the maximum keep time and the maximum number of versions to retain. The Delete Policy for volumes can be configured to allow indefinite retention of at least one most recent file version of a current file.
- ♦ To retrieve file versions using the NSS File Version Utility, users must download and install the NSS File Version utility on their Windows 2000/XP/ Vista or SLED 10/ SLED 10 SP1 workstations.

NOTE: If the filename or directory name contains non-english characters, NSS File Version Utility fails to restore versions to a different location on Windows 2000/ XP/ Vista workstation. Instead you can use NetStorage to restore files to a different location.

- ♦ To retrieve file versions using Novell NetStorage, users must have a workstation with a compatible Web browser and a network or Internet connection to the Novell NetStorage server.

4.10 Guidelines for Availability of File Versions

After your system meets the [Prerequisites and Guidelines for Retrieving File Versions](#), there are several reasons that previous versions of a current, renamed, or deleted file might not exist in the archive database:

File Versions Are No Longer Supported for a Volume

If you suspend or delete a versioning job for a volume, its scheduled jobs do not run and no new versions are saved to the archive. For suspended jobs, users can retrieve available file versions from the archive. For suspended jobs, the job's Delete Policy continues to be applied to existing versions in the archive. Eventually, the file versions meet the criteria for deletion defined in the job's Delete Policy.

If Keep Current Copy is enabled for the job's Delete Policy, at least one file version of a file remains in the database if its source file is current on the source volume, even if the Maximum Keep Time elapses.

If Keep Current Copy is disabled, the archive deletes the file version when the Maximum Keep Time elapses.

A File's Lifetime Does Not Span the End of an Epoch

A file is versioned only if it exists at the end of an epoch. If scheduled versioning processes are paused or delayed for a period of time that exceed the lifetime of a given file, its file version might not be captured, even if the file exists for a length of time that exceeds the epoch. Although the frequency of the epoch usually is less than the typical lifetime of your files, there might be occasions when the versioning time is extended. For example, if the time it takes to save versions of all the files eligible for versioning overlaps one or more scheduled epochs, some scheduled epochs might be skipped until the current process ends and the next scheduled start time occurs.

To avoid this problem:

- ♦ Schedule a job's epoch for a period of time shorter than a typical file's lifetime for the volume, while allowing enough time to capture file versions.
- ♦ Enable the Pool Snapshot option for capturing file versions.
- ♦ Minimize the use and length of job pauses.

A File Is Open When Versioning Occurs

Novell Archive and Version Services 2.0 leverages NSS pool snapshot technology to save point-in-time versions of all files, including open ones. If you disable the Pool Snapshot option, the files that are eligible for versioning are copied directly from the source volume. Exclusively open files cannot be versioned and data might be inconsistent.

To avoid this problem, make sure to enable the Pool Snapshot option for the archive server.

A File Is Deleted Before A Version Can Be Saved

If the Pool Snapshot option is disabled for your archive server, file versions are saved from the volume to the archive, which takes time. If a file is created during an epoch and exists at the end of an epoch, but a user deletes it before a version can be copied to the archive database, that file version is not saved to the archive database. If the Pool Snapshot option is enabled, versions of all files, even exclusively open ones, are captured by the snapshot and a point-in-time version of each file can be saved to the archive server.

To avoid this problem, make sure to enable the Pool Snapshot option for the archive server.

The Age of a File Version Exceeds Retention Thresholds

If your Delete Policy is excessive, file versions might exceed retention thresholds too quickly to be of use to your users. Typically, the file's archived versions are deleted automatically from the archive database as they exceed the Maximum Time to Keep Versions or the Maximum Number of Versions.

To avoid this problem, make sure your Delete Policy meets the needs of your users.

4.11 What's Next

After your deployment plan and resources meet the prerequisites and guidelines, you are ready to install or upgrade Archive and Version Services. See the following:

- ♦ [Chapter 5, "Setting Up Archive and Version Services," on page 41](#)

5 Setting Up Archive and Version Services

This section describes how to install and configure Novell Archive and Version Services 2.1 on OES 11 server.

- ♦ [Section 5.1, “Before you begin,” on page 41](#)
- ♦ [Section 5.2, “Installing and Configuring an Archive Server,” on page 41](#)
- ♦ [Section 5.3, “Upgrade,” on page 46](#)
- ♦ [Section 5.4, “Reinstall,” on page 48](#)
- ♦ [Section 5.5, “What’s Next,” on page 48](#)

5.1 Before you begin

Before installing Novell Archive and Version Service 2.1, review the following:

- 1 Plan your Archive server implementation.
For more information, see
 - ♦ [Chapter 3, “Planning for Archive and Version Services,” on page 21](#)
 - ♦ [Chapter 12, “Security Considerations for Archive and Version Services,” on page 99](#)
- 2 Ensure that the Archive server implementation satisfies the [“Prerequisites and Guidelines” on page 33](#) for Novell Archive and Version Services.
- 3 Continue with [Section 5.2, “Installing and Configuring an Archive Server,” on page 41](#)

IMPORTANT: For security recommendations on OES 11 server, you must install the server patch to run the Archive and Version Services as a non-root user. For more information, see the [“Server Patch” on page 44](#).

For information on installing OES 11 services, see the [OES 11: Installation Guide](#).

5.2 Installing and Configuring an Archive Server

This section discusses the following tasks:

- ♦ [Section 5.2.1, “Unsupported Service Combinations,” on page 42](#)
- ♦ [Section 5.2.2, “YaST Install,” on page 42](#)
- ♦ [Section 5.2.3, “Creating an NSS Pool,” on page 44](#)
- ♦ [Section 5.2.4, “Creating NSS Volumes,” on page 44](#)

- ♦ [Section 5.2.5, “Configuring Jobs using iManager,” on page 45](#)
- ♦ [Section 5.2.6, “Restoring Versions,” on page 45](#)

5.2.1 Unsupported Service Combinations

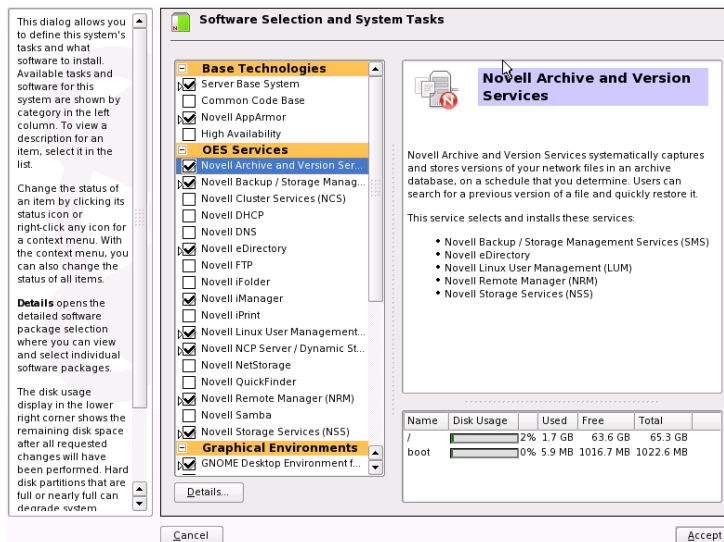
IMPORTANT: Do not install any of the following service combinations on the same server as the Novell Archive and Version Services.

Although not all of the combinations will cause pattern conflict warnings, Novell does not support any of the combinations shown.

- ♦ Novell Domain Services for Windows
- ♦ Xen Virtual Machine Host Server

5.2.2 YaST Install

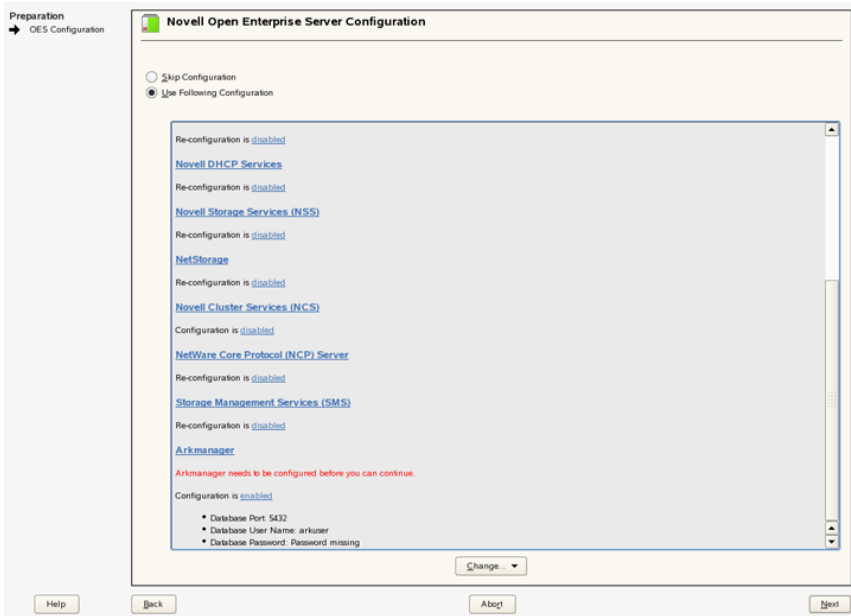
- 1 During installation of OES Services, in the Software Selection page, select Novell Archive and Version Services option, the dependent components are selected along with the service. You have to additionally select the following dependencies:
 - ♦ **Novell iManager:** Allows to configure server, defaults, and job settings and to manage jobs.
 - ♦ **Novell NetStorage:** Provides a Web based interface to users for file version retrieval and restoration.



Click *Accept* and follow the on-screen instructions.

- 2 On the Novell Open Enterprise Server Configuration screen, you can do one of the following:
 - ♦ Click *Next* to use the default Arkmanager configuration and continue to [Step 3](#). If you use the default Arkmanager configuration, Archive service starts its own instance of PostgreSQL database.
 - ♦ If you need to configure Archive and Version Services with an existing PostgreSQL database, select *Arkmanager* to change the credentials of database and continue to [Step 2a](#).

NOTE: The existing PostgreSQL database needs to be running. Archive server creates its own instance of database, if it is not running.



2a (Optional) On selecting Arkmanager, a Archive Versioning Server Configuration screen is displayed. Enter the inputs for the fields, refer [Table 5-1](#) or the help on the left pane of the screen for details on the fields.

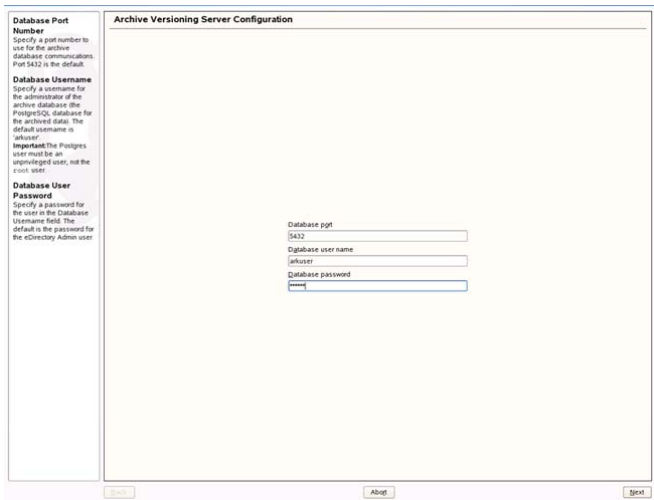


Table 5-1 Default Configuration File Parameter Description

Parameter	Description
Database Port Number	Specify the port number to use for the archive database communications. Port 5432 is the default.
Database Username	Specify a username for the administrator of the archive database (the PostgreSQL database for the archived data). IMPORTANT: The PostgreSQL user must be an unprivileged user, not the <code>root</code> user.
Database User Password	Specify the password for the user specified in the <i>Database Username</i> field.

The database runs with the credentials configured by you.

3 Click *Next*

Follow the on-screen instructions to complete the Archive server configuration.

4 The Archive server is now configured with the same eDirectory credentials you have logged in. You must create an [NSS Pool](#) and [NSS Volumes](#) to store the archive data.

Server Patch

To run the Archive and Version Services as non-root user, you must install the server patch.

1 Install the Archive Server Patch # 15269. On installation of the patch, you need to reconfigure the Archive and Version Services.

IMPORTANT: If you have already configured the Archive server, you can retain the existing configuration and user data. For more information, see the [TID 3858638 \(http://www.novell.com/support/search.do?cmd=displayKC&docType=kc&externalId=3858638&sliceId=SAL_Public&dialogID=54020695&stateId=0%200%2054030232\)](http://www.novell.com/support/search.do?cmd=displayKC&docType=kc&externalId=3858638&sliceId=SAL_Public&dialogID=54020695&stateId=0%200%2054030232) at Novell Technical Support website.

5.2.3 Creating an NSS Pool

You must create an NSS pool on the server where you plan to store the archive data. For information, see Section 16.1, Creating a Pool in the [OES 11: NSS File System Administration Guide for Linux](#).

5.2.4 Creating NSS Volumes

You must create two encrypted or non-encrypted NSS (Primary and Archive) volumes on same or different [NSS](#) pool.

The Primary volume contains the data to be versioned and Archive volume contains the versions or the archived data. One versioning job can be defined for each Primary volume.

To create an non-encrypted NSS volume through iManager, see Section 19.3, Creating and Configuring Unencrypted NSS Volumes in the [OES 11: NSS File System Administration Guide for Linux](#).

If your implementation requires an encrypted NSS volume to store the file versions, use the NSS Management Utility (NSSMU) to create the encrypted volume. The iManager Storage plug-in does not provide an encryption option. For information, see Section 20.3 Creating an Encrypted Volume in the [OES 11: NSS File System Administration Guide for Linux](#).

To archive file versions, you must define versioning jobs to run on the archive server. Continue with the next section for [configuring jobs](#) using iManager.

5.2.5 Configuring Jobs using iManager

A job specifies the property settings to use when archiving file versions for a specified volume that resides in the same eDirectory tree as the archive server.

You must establish a schedule for each versioning job that meets users' requirements for file versioning, given limited storage and bandwidth resources. Versioning occurs for eligible files at scheduled intervals, called epochs. We recommend you to schedule 30-minute epochs for the volume.

You can only define one job for a given Primary volume. If you attempt to define multiple jobs for a volume on the same or different server, Archive and Version Services does not run as designed and data integrity in the archive database is compromised.

The iManager utility allows you to configure Archive volume and create and manage jobs. For information on configuring jobs using iManager, see [Chapter 7, "Configuring Jobs in iManager," on page 53](#).

5.2.6 Restoring Versions

After configuring a job, you can access file versions using NSS File Version Utility or Web based Archive function using Novell NetStorage. For more information on restoring files, see [OES 11: Novell Archive and Version Services 2.1 User Guide](#).

You must LUM-enable the user who needs to restore the file versions using NSS File Version Utility. To do this, Click the *Linux User Management > Enable Users for Linux* task.

In the example, we have used the name of Primary server as *Arcprime-server*. The context of the server is *finance*. The name of the user is *Carlos* and group is *Archivegp*.

- 1 In the *Select Users* window, select *Carlos* and click *Next*.
- 2 Every Linux user must belong to a primary group. In the *Select Primary Group* window, select *novlxtier* object and click *Next*.

IMPORTANT: A user must be a part of the *novlxtier* group to use the NSS File Version Utility for restoring versions.

- 3 In the *Select Workstations* window, select *UNIX WorkStation-Arcprime-server.finance* and click *Next*.

The user *Carlos* is LUM-enabled and his primary group is *novlxtier*. *Carlos* is also given access to *UNIX WorkStation-Arcprime-server.finance* workstation.

NOTE: In addition to *novlxtier* group, user can be added to another groups using iManager. To add a user to *Archivegp* group, click *Users>Modify User* task. Select the user *Carlos* and click OK. Select the tab *Group Memberships* and add *Archivegp* group.

The user *Carlos* is now a member of *novlxtier* and *Archivegp* groups.

For more information on LUM, refer the [OES 11: Novell Linux User Management Administration Guide](#).

5.3 Upgrade

To upgrade Archive and Version Services to the OES 11 server, you must backup archive database on the OES 2 server and restore the Archive database to the OES 11 server.

NOTE: The archive database (PostgreSQL) needs to be backed up and restored because PostgreSQL does not support upgrading from version 8.1 to 8.3.

- [Section 5.3.1, “Before You Upgrade,” on page 46](#)
- [Section 5.3.2, “Post-Upgrade,” on page 46](#)
- [Section 5.3.3, “Caveats on Upgrading from OES 2 to OES 11,” on page 47](#)

5.3.1 Before You Upgrade

Before you upgrade to OES 11, run the script to backup the Archive database on the OES 2 server. The user must have execute permissions to run the script. The script backs up the archive database to the `arkdatabackup.sql` file.

- 1 Save the following script to a file, then run the file on the terminal:

```
#!/bin/bash
set -e
echo
echo "Backing up Archive Versioning archive_database"
echo

ARKUSER=arkuser
DATABASE=archive_database
DATA_PATH=`cat /etc/opt/novell/arkmanager/conf/arkdatadir.conf`
BACKUPFILE=$DATA_PATH/arkdatabackup.sql

echo "Provide password for $ARKUSER"
su $ARKUSER -c "pg_dump --clean $DATABASE > $BACKUPFILE"
```

You are prompted for the `arkuser` password.

NOTE: If the `arkdatabackup.sql` file is empty, repeat [Step 1](#) to generate the contents.

5.3.2 Post-Upgrade

On upgrading to the OES 11 server, you must restore the Archive database for Archive and Version Service to be available. The user must have execute permissions to run the script. The script restores the archive database from the `arkdatabackup.sql` file to the OES 11 server.

- 1 Configure Archive Version and Services using the same credentials as used on the OES 2 server.

```
#!/bin/bash
set -e
echo
echo "Restoring Archive Versioning archive_database"
echo

ARKUSER=arkuser
DATABASE=archive_database
DATA_PATH=`cat /etc/opt/novell/arkmanager/conf/arkdatadir.conf`
BACKUPFILE=$DATA_PATH/arkdatabackup.sql
LOGFILE=$DATA_PATH/restorelog.txt
PORT=543
HOST=localhost

if [ -z DATA_PATH ]
then
    echo "Unable to locate $DATABASE location"
    exit 1;
else
    echo "Provide password for $ARKUSER"
    su $ARKUSER -c "psql -f $BACKUPFILE -d $DATABASE -h $HOST -p $PORT"
    >>$LOGFILE 2>&1
fi
```

You are prompted for the arkuser password.

On successfully restoring the archive database, the file versions are available on the OES 11 server.

5.3.3 Caveats on Upgrading from OES 2 to OES 11

On upgrading the Archive and Version Services from the OES 2 server to OES 11 server, Archive server may fail to come up. The failure can be caused due to the following reasons:

- ♦ The group id (gid) of both the novlxtier group (group in eDirectory tree) and the local user group (group on the local machine) are same.
- ♦ Admin user is not listed as the part of the novlxtier group.

To resolve the group ids conflicts, perform the following steps, then restart the Archive server.

- 1 To verify the groups of the admin, enter

```
id admin
```

For example, the output of the command is

```
uid=600(admin) gid=600(admingroup) groups=600(admingroup) ,8(www)
,101(messagebus)
```

The novlxtier group is not listed.

- 2 The novlxtier gid must be 81. To verify if any object is using the same gid (81) of the novlxtier group, enter

```
getent group | grep 81
```

- 3 To change the gid of the novlxtier group to 81, enter

```
namgroupmod -a adminname -w password -g 81 -o cn=novlxtier,servercontext
```

On execution of the above command the gid of novlxtier group changes to 81.

For more information on the namgroupmod utility, see “[namgroupmod](#)” in the *OES 11: Novell Linux User Management Administration Guide*.

- 4 Changing the gid for novlxtier group may cause ownership issues on the files. To resolve the file ownership issues, execute the `nssid.sh` script.

```
/opt/novell/oes-install/nssid.sh
```

For more information, see “[The OES 1 Solution: The nssid.sh Script](#)” in the *OES 11: Planning and Implementation Guide*.

- 5 To ensure that the admin user is part of the novlxtier group, enter

```
id admin
```

For example, the output of the command is

```
uid=600(admin) gid=600(admingroup) groups=600(admingroup) ,8(www)
,81(novlxtier)
```

The admin is part of the novlxtier group.

- 6 To restart the Archive server, enter

```
rcnovell-ark restart
```

5.4 Reinstall

Reinstall the Archive and Version Services using YaST.

- 1 Stop the Archive server.

```
rcnovell-ark stop
```

- 2 Open YaST, go to *Open Enterprise Server>OES Install and Configuration* screen.

- 3 Novell Archive and Version Services is already selected in the Software Selection page. Click *Accept*.

- 4 On the Novell Open Enterprise Server Configuration screen, under *Arkmanager*, select the *disabled* option. To change the database credentials complete the following steps:

- 4a Click *ArkManager*, you are prompted for admin password.

- 4b On the Archive Versioning Server Configuration screen, change the database credentials and click *Next*.

- 5 Click *Finish*. The Archive server is configured with new database credentials.

- 6 Start the Archive server

```
rcnovell-ark start
```

The Archive server runs with new credentials.

5.5 What's Next

Proceed to one of the following sections based on the task that you want to perform:

- ♦ [Chapter 11, “Migrating Novell Archive and Version Services to OES 11,”](#) on page 97
- ♦ [Chapter 8, “Managing Jobs,”](#) on page 71
- ♦ [Chapter 6, “Managing the Archive Server,”](#) on page 49

6 Managing the Archive Server

This section discusses how to manage your Novell Archive and Version Services 2.1 on OES 11 server.

- ♦ [Section 6.1, “Starting ArkManager,” on page 49](#)
- ♦ [Section 6.2, “Stopping ArkManager,” on page 49](#)
- ♦ [Section 6.3, “Updating the Archive Server Configuration,” on page 50](#)
- ♦ [Section 6.4, “Vacuuming the Archive Database,” on page 51](#)
- ♦ [Section 6.5, “Backing Up the Archive Database,” on page 51](#)
- ♦ [Section 6.6, “Backing Up the Archive Data,” on page 51](#)

6.1 Starting ArkManager

Starting ArkManager on a Server

To start ArkManager, enter the following command at the server console command prompt:

```
rcnovell-ark start
```

If the service was running earlier, Archive server starts automatically whenever you restart the server.

6.2 Stopping ArkManager

Stopping ArkManager on a Server

To stop ArkManager, enter the following command at the server console command prompt:

```
rcnovell-ark stop
```

After typing `arkstop`, make sure ArkManager is shutting down cleanly before running it again. You can check this in two ways:

- ♦ In iManager if the *Display Log* option is enabled for the server, the log messages are displayed at *ArchiveVersioning > Archive Jobs > Log* tab.
- ♦ On the archive server’s console screen, use the `rcnovell-ark status` command to check that no ArkManager process is running.

It might take a while for ArkManager to shut down because it needs to wait for its current file archiving process to finish. Under normal circumstances, ArkManager should eventually shut down cleanly.

6.3 Updating the Archive Server Configuration

Archive server authenticates the credentials using eDirectory and PostgreSQL. If you have modified the eDirectory or PostgreSQL credentials, Archive server fails with an authentication error, unless you have updated the modified details in the Archive server.

- ♦ [Section 6.3.1, “Updating the Archive Server Configuration with Modified eDirectory Credentials,” on page 50](#)
- ♦ [Section 6.3.2, “Updating the Archive Server Configuration with Modified PostgreSQL Credentials,” on page 50](#)

6.3.1 Updating the Archive Server Configuration with Modified eDirectory Credentials

If you have modified the eDirectory credentials, Archive server must be updated to reflect the changes.

- 1 Stop the Archive server

```
rcnovell-ark stop
```

- 2 Go to /opt/novell/arkmanager/bin folder and run the script arkReconfig.sh:

```
sh arkReconfig.sh --edir-admin=<username>
```

where <username> is the same or modified eDirectory admin user.

You will be prompted to enter the password for the eDirectory admin.

NOTE: This command changes both the username and password.

- 3 To run with modified credentials, you need to start the Archive server:

```
rcnovell-ark start
```

The Archive server is updated with the modified user credentials.

6.3.2 Updating the Archive Server Configuration with Modified PostgreSQL Credentials

If you have modified the PostgreSQL configuration details, Archive server must be updated to reflect the changes.

- 1 Stop the Archive server

```
rcnovell-ark stop
```

- 2 Go to /opt/novell/arkmanager/bin folder.

- 2a If you have modified the PostgreSQL credentials, run the script arkDBReconfig.sh

```
sh arkDBReconfig.sh --db-user=<username>
```

where <username> is the database user.

You will be prompted to enter the new password.

NOTE: This command changes both the username and password.

- 2b If you have modified the database port number, run the script ark_change.sh

```
sh ark_change.sh --ark-port=<portnumber>
```

where <portnumber> is the modified database port number.

- 3 To run with new PostgreSQL credentials, you need to start the Archive server:

```
rcnovell-ark start
```

The Archive server is updated with the modified user credentials.

6.4 Vacuuming the Archive Database

The Archive server stores the metadata information of the versions of the files in the PostgreSQL database. The metadata information is stored in multiple tables that are associated with specific jobs. Based on the policies defined for each job, multiple records are created for each of the versioned files.

The archive database (PostgreSQL) requires the VACUUM command to scan the tables, marking that are no longer needed as free space so that they can be overwritten by newly inserted or updated data. Based on the size of your database, you can vary the frequency of running the VACUUM command.

For each job, ArkManager creates multiple tables such as file table (arkdataxxxxx_f), user table (arkdataxxxxx_u), history table (arkdataxxxxx_h) and metadata table (arkdataxxxxx_m). Since the file table contains the maximum entries, this needs to be vacuumed periodically. To run the VACCUM command on the file table, execute the following:

- 1 At the server console command prompt, run the following command:

```
vacuumdb -U arkuser -d archive_database -z -t arkdataxxxxx_f -v -e
```

Vacuum on arkdataxxxxx_f table increases the performance of the database.

- 2 Specify the arkuser password when prompted.

Repeat this procedure for every file table in the database.

6.5 Backing Up the Archive Database

You should periodically back up the archive database. The frequency depends on the critical nature of the versions you archive.

WARNING: To ensure data integrity, shut down ArkManager before and during the backup. For information, see [Section 6.2, “Stopping ArkManager,” on page 49](#).

For backup and restore operations, see [PostgreSQL 8.1 Manual \(http://www.postgresql.org/docs/8.1/static/backup.html\)](http://www.postgresql.org/docs/8.1/static/backup.html).

6.6 Backing Up the Archive Data

You should periodically back up the archive data. The frequency depends on the critical nature of the versions you archive.

WARNING: To ensure data integrity, shut down ArkManager and PostgreSQL before and during the backup. For information, see [Section 6.2, “Stopping ArkManager,” on page 49](#).

To back up the archive versions for a particular job, copy the directories to a different volume. We recommend that the destination volume where you store the backup copy be on a different drive and pool than your archive volume, although it is not mandatory. You can also use your standard backup tools and procedures to make the directories part of your scheduled backup.

The general data path is

```
/media/nss/ARK/archive/arkDataxxxxxxx/
```

where *ark* is the Archive volume, *archive* is the archive data directory, and *xxxxxxx* is a 6-digit random number assigned to the job.

If the archive path is in a directory in the volume, simply copy all files in the directory.

7 Configuring Jobs in iManager

To make recent file versions available to your users, you must set up one or more Novell Archive and Version Services servers. This section discusses how to configure versioning jobs, using the Archive Versioning plug-in for Novell iManager.

- ♦ [Section 7.1, “Accessing the Archive Versioning Plug-In in iManager,” on page 53](#)
- ♦ [Section 7.2, “Configuring Archive Volume,” on page 54](#)
- ♦ [Section 7.3, “Configuring or Viewing Archive Server Properties,” on page 54](#)
- ♦ [Section 7.4, “Configuring Default Job Settings,” on page 56](#)
- ♦ [Section 7.5, “Configuring Job Properties,” on page 60](#)
- ♦ [Section 7.6, “Creating a Job,” on page 65](#)
- ♦ [Section 7.7, “Editing a Job’s Properties,” on page 67](#)
- ♦ [Section 7.8, “What’s Next,” on page 69](#)

7.1 Accessing the Archive Versioning Plug-In in iManager

- 1 Open your Web browser to the following URL:

`https://svrname.example.com/nps/iManager.html`

Replace *svrname.example.com* with the actual DNS name or IP address (for example, 192.168.1.255) of the server where iManager is running.

IMPORTANT: The URL path is case sensitive.

For information, see Accessing iManager chapter in the [Novell iManager 2.7.4 Administration Guide](#).

- 2 On the iManager Login page, log in to the eDirectory tree where the archive server you want to manage resides.
- 3 In the left navigator, expand Archive Versioning to show its tasks.



- 4 Click *Archive Jobs* or *Archive Server Properties*, depending on what task you want to perform.
- 5 Select the archive server you want to manage in the tree where you are logged in to iManager.
For information, see [Section 7.2, “Configuring Archive Volume,” on page 54](#).

7.2 Configuring Archive Volume

- 1 Click *Archive Versioning > Archive Server Properties* to open the Archive Server Properties window in the main panel.
- 2 In the *Server* field, specify or browse the name of the eDirectory™ context for the archive server you want to manage, then wait for the page to refresh.

Archive Server Properties

Server: Archive-server.novell

Properties Default Job Settings

Specify the volume and path on the archive server where you want to store file versions. Select Display Archive Log to also send log entries to the archive server's console. The admin users who manage Archive Versioning and PostgreSQL (MySQL on NetWare) archive database are specified in the /etc/opt/novell/arkmanager/conf/arkConfig.xml (sys:\ArkManager\arkConfig.xml on NetWare) file.

Volume Path: ARK:/archive (ex. ARK:/archive)

☒ Display server Archive Log entries

Authentication

eDirectory:

User Name: cn=admin.o=novell

Database:

User Name: arkuser

Port: 5432

OK Cancel Apply

- 3 In the *Volume Path* field, specify the volume and path on the archive server to store the versions of files.
The *Authentication* field displays, eDirectory username, PostgreSQL username and port.
- 4 Click OK.

7.3 Configuring or Viewing Archive Server Properties

The Archive Server Properties page displays the basic management information that applies to all jobs controlled by the selected archive server. The general server properties specify the storage location where the archive data (file versions) reside and whether to display the log using iManager at *Archive Versioning > Archive Jobs > Log* tab.

- ♦ [Section 7.3.1, “Accessing the Server Properties Page,” on page 54](#)
- ♦ [Section 7.3.2, “Setting Archive Server Properties,” on page 55](#)
- ♦ [Section 7.3.3, “Viewing Authentication Properties,” on page 55](#)

7.3.1 Accessing the Server Properties Page

To manage the archive server's properties:

- 1 In iManager, expand *Archive Versioning*, select *Archive Server Properties*, then click *Properties*.
- 2 Select the archive server you want to manage, then wait for the page to refresh.
For information, see [Section 7.2, “Configuring Archive Volume,” on page 54](#).

7.3.2 Setting Archive Server Properties

- 1 On the Server Properties page, set the following server properties for the archive server:

Property	Description
<i>Volume Path</i>	<p>Specifies the path to the volume on the archive server where the file versions are stored for retrieval by users. The directory can be the root of the volume or any other directory in the volume. For example, on OES 11, a Novell Storage Services volume path might be <i>Ark:/archive</i>. If you are setting up a clustered solution for the archive server using Novell Cluster Services, make sure to specify the volume and path to the virtual storage location.</p> <p>NOTE: The Archive volume name must always be in uppercase.</p>
<i>Display server Archive Log entries</i>	<p>By default, the archive server records the error, warning, and normal messages for all of its jobs in the Archive Log available at <code>/var/opt/novell/arkManager/logs</code> file.</p> <p>In iManager if the <i>Display server Archive Log entries</i> option is selected, the log messages are displayed at <i>Archive Versioning > Archive Jobs > Log</i> tab. This is the default setting</p> <p>If <i>Display Archive Log Entries</i> is deselected, the archive server does not print messages to the <i>Log</i> tab, but the messages are recorded in the log file.</p> <p>For information, see Section 8.5, “Viewing the Archive Log,” on page 76.</p>

- 2 When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

7.3.3 Viewing Authentication Properties

- 1 On the Server Properties page, view the authentication information for Novell eDirectory and PostgreSQL that you configured using YaST.

Property	Description
<i>eDirectory</i>	<p>User Name Specifies the Novell eDirectory common name of the administrator user who has the appropriate rights to the original data location and to the archive data location. For example,</p> <p><code>admin.servercontext</code></p> <p>The archive server administrator user must have supervisor rights to all servers being accessed by the selected archive server. The user must be in the same tree as the archive server and the primary server.</p>
<i>Database</i>	<p>Specifies authentication information about the PostgreSQL database for the Archive and Version Services server.</p> <p>User Name Specifies the administrator user of the PostgreSQL database.</p> <p>Port Specifies the port number to use for the archive database communications. Port 5432 is the default.</p>

- 2 When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

7.4 Configuring Default Job Settings

Default job settings are property settings that can optionally be used in any individual job on the server. Go to an individual job to configure it to use defaults.

- [Section 7.4.1, “Accessing the Default Job Settings Page,” on page 56](#)
- [Section 7.4.2, “Setting Default Job Information,” on page 56](#)
- [Section 7.4.3, “Setting Default Source Server Information,” on page 56](#)
- [Section 7.4.4, “Setting Default Run Schedule Information,” on page 57](#)
- [Section 7.4.5, “Setting Default Delete Policy Information,” on page 58](#)
- [Section 7.4.6, “Applying Default Job Settings,” on page 59](#)

7.4.1 Accessing the Default Job Settings Page

To manage the archive server’s default job settings:

- 1 In iManager, expand *Archive Versioning*, then click *Archive Server Properties*.
- 2 Select the archive server you want to manage, then wait for the page to refresh.
For information, see [Section 7.2, “Configuring Archive Volume,” on page 54](#).
- 3 Select the *Default Job Settings* tab.

7.4.2 Setting Default Job Information

- 1 On the Default Job Settings page, specify the following default job information property.

Property	Description
<i>Copy All Files the First Time the Job Is Run</i>	If <i>Copy All Files</i> is selected, the job saves versions of all eligible files to the archive database the first time the job is run. Thereafter, the job saves file versions only for files that changed. If <i>Copy All Files</i> is deselected, the job saves versions only of eligible files that were changed during the epoch, such as files that were added, modified, renamed, deleted, or where trustee information changed.

- 2 When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

NOTE: On configuring the job first time, Copy All Files may get deselected on refresh. Select again before clicking Apply.

7.4.3 Setting Default Source Server Information

- 1 On the Default Job Settings page, specify the following default source server information properties.

Property	Description
<i>Server</i>	<p>The host name of the source server where the data to be versioned is located. For example, <code>servername.context</code>. The source server is in the same tree as the Archive and Version Services server.</p> <p>If you specify a new source server for an existing job, the archived data becomes associated with the source volume on the new source server. Typically, you change the source server only when you have renamed the it.</p>
<i>Snapshot Pool</i>	<p>The name of the destination pool where snapshots of the source volume are created and temporarily maintained at the end of an epoch until point-in-time file versions can be saved to the archive database. For example, <code>pusers_sl</code> or <code>pdata_sl</code>.</p> <p>If no snapshot pool name is specified, or if the snapshot cannot be created for any reason, the versioning process copies files directly from the source volume.</p>
<i>Free Space ID</i>	ID of the free space object to be used for storing snapshot data. For example, <code>/dev/hda</code> .
<i>Sectors</i>	Specifies the number of sectors on free space to be used for storing snapshot data. The value needs to be an integer, for example, 409600.
<i>Filter</i>	Set filters to specify the types of files and directories to be versioned. For more information, refer to the online help.

- 2 When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

7.4.4 Setting Default Run Schedule Information

- 1 On the Default Settings page, specify when to start the job and the frequency for running the job. To set the frequency, you must specify only one of three scheduling options:

Property	Description
<i>No Run Schedule</i>	Select <i>No Run Schedule</i> to indicate that there is no default schedule available for the server.
<i>Every</i>	<p>Select <i>Every</i> to enable a <i>Scheduled Interval</i>. Specify the elapsed time between the beginning of versioning processes for a job. In the Units drop-down list, select seconds, minutes, hours, or days. For example, 45 seconds, 1 minute, 15 minutes, 2 hours, or 12 hours.</p> <p>If the versioning process exceeds the time specified as the interval, the overlapping scheduled job is skipped. No file versions are saved for skipped job runs. After the version process completes, the job runs at its next scheduled interval. If you observe that the job skips some versioning intervals, you can increase the interval between versions, or you can reduce the amount of data to be versioned by setting Filter properties in the Source Server Information section using iManager.</p>
<i>Time</i>	<p>Select <i>Time</i>, specify the start time when the job's version process begins, then specify one or more days of the week to run the job.</p> <ul style="list-style-type: none"> ♦ Time: Click the double-arrows to navigate to the time by hours or click the single-arrow to navigate by 15-minute increments. For example, 00:00 AM. ♦ Days of the Week: Select the <i>Days of the Week</i> check box to select all days for a daily job run, or select only the check boxes next to one or more days of the week you want the job to run. Choices include <i>Monday</i>, <i>Tuesday</i>, <i>Wednesday</i>, <i>Thursday</i>, <i>Friday</i>, <i>Saturday</i>, <i>Sunday</i>, and <i>All</i>.

- 2 When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

7.4.5 Setting Default Delete Policy Information

The delete policy determines the retention of file versions by age or by number of versions. If a delete policy is set, the job runs a process to delete file versions according to its own Delete Schedule. The delete process is not related to the job's Run Schedule, which determines when file versions are saved from the source volume. The job's delete policy runs if the job is in a Running, Scheduled, or Stopped state. The job's delete policy does not run if a job is in the Clean Up Users, Not Configured, or Deleted state.

IMPORTANT: The Delete Schedule operates separately from the Run Schedule.

- 1 On the Default Settings page, specify one of the following options:

Property	Description
<i>No Delete Policy</i>	Select this option if you want to retain file versions indefinitely.
<i>Define Job Delete Policy</i>	Select this option if you want to configure a default delete policy, then specify the schedule for deleting file versions, if they are eligible for deletion. Set the <i>Interval</i> and <i>Maximum Keep</i> settings.

- 2 If you selected *Define Job Delete Policy*, complete the following information:

Property	Description
<i>Interval</i>	<p>This setting represents the amount of time to wait from the time a delete-versions process ends until another delete-versions process begins. If a value is not specified, 24 hours is the default interval. How long the delete process takes depends on the number of files stored in the archive database.</p> <p>For example, suppose you set the <i>Delete Schedule Interval</i> to 1 hour. When you save a job, ArkManager starts the interval timer. After 1 hour elapses, the job kicks off its delete-versions process, resets the interval timer, and pauses the timer. When the delete process ends, the interval timer begins. The delete intervals repeat in this manner until the job is stopped.</p>
<i>Maximum Keep</i>	<p>Specifies the maximum number of versions of each file to keep in the archive and how long to keep file versions. At least one of the values must be non-zero. If you set both the <i>Maximum Keep Versions</i> and <i>Maximum Keep Time</i> to zero values, the Delete Policy function does not run.</p> <p>NOTE: Whenever the job is run, it checks for the values specified in the <i>Time</i> and <i>Versions</i> fields to delete the job. The job is deleted depending on the condition that is satisfied, by the values specified in either the <i>Time</i> or <i>Versions</i> field.</p> <ul style="list-style-type: none"> ♦ Time: The maximum time that a file version is maintained in the archive. Specify the time in whole numbers; then in the <i>Units</i> drop-down list, select seconds, minutes, hours, or days. ♦ Versions: Specify the maximum number of versions of each file to keep in the archive. As the number of versions exceeds this integer value, the oldest version is deleted. ♦ Keep Current Copy: By default, the latest file version of current files remains in the database, even if the <i>Maximum Keep Time</i> elapses. <p>If <i>Keep Current Copy</i> is selected, the archive keeps the latest file version as long as its source file is current on the source volume, even if the <i>Maximum Keep Time</i> elapses. After the user deletes the current source file, the deletion is noted when the versioning process runs. If the file version's age is within the <i>Maximum Keep Time</i>, the archive database retains the file version of the deleted file; otherwise, the file version is deleted.</p> <p>If <i>Keep Current Copy</i> is deselected, the archive deletes the file version when the <i>Maximum Keep Time</i> elapses.</p>

- 3 When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

7.4.6 Applying Default Job Settings

Each job can optionally use none, one, or more of the default settings. Each job's usage of defaults is independent of other jobs' usage.

It is not necessary to stop jobs that use defaults while you modify, add, or remove default settings. Make your changes, then click *OK* or *Apply* to save them. The following table describes special circumstances for how the changes take effect.

For more information on the state of the Job, refer [Job Status information](#).

Table 7-1 How Changes to Default Job Settings Take Effect

Change Made to Default Settings	State of the Job When You Save Default Job Settings	Additional Actions in Jobs that Use Defaults
Modifying an existing default Delete Policy	Running, Scheduled, or Stopped	If a delete-versions process is in progress, the run is completed with the old delete policy. The delete policy applies for the next delete-versions.
Modify any existing default setting except Delete Policy	Running	The job runs are completed with the old setting.
	Clean Up Users	The Clean Up Users run is stopped gracefully, and the job is placed in the Stopped state. The list might be only partially cleaned up.
	Not Configured	The job remains in the Not Configured state until you verify the individual job settings and start or schedule the job.
Add a default setting	Any	To take advantage of the new defaults, you must modify settings in individual jobs to use them.
Remove default settings	Running	The job run is completed with the old settings, then the job is placed in the Not Configured state. Go to the individual jobs to specify values for the missing settings.
	Scheduled or Stopped	Jobs are placed in the Not Configured state. Go to the individual jobs to specify values for the missing settings.
	Clean Up Users	The Clean Up Users run is stopped gracefully, and the job is placed in the Stopped state. The users list might be only partially cleaned up. The job is then placed in the Not Configured state. Go to the individual job to specify values for the missing settings.
	Not Configured	The job remains in the Not Configured state until you verify the individual job settings and start or schedule the job.

7.5 Configuring Job Properties

Use the following guidelines for setting job properties on the [Create New Job](#) page and [Edit Job](#) page.

- ♦ [Section 7.5.1, “Accessing the Job Settings Page,”](#) on page 60
- ♦ [Section 7.5.2, “Setting Job Information,”](#) on page 61
- ♦ [Section 7.5.3, “Setting Source Server Information,”](#) on page 61
- ♦ [Section 7.5.4, “Setting Run Schedule Information,”](#) on page 62
- ♦ [Section 7.5.5, “Setting Delete Policy Information,”](#) on page 63

7.5.1 Accessing the Job Settings Page

To manage the archive server’s default job settings:

- 1 In iManager, expand *Archive Versioning*, then click *Archive Server Properties*.
- 2 Select the archive server you want to manage, then wait for the page to refresh.

For information, see [Section 7.2, “Configuring Archive Volume,” on page 54.](#)

- 3 Select the *Job Settings* tab.

7.5.2 Setting Job Information

- 1 On the Job Settings page, configure the following job information properties:

Property	Description
<i>Name</i>	The administrator-specified unique job name. For example, <code>svr1_users</code> or <code>svr1_data</code> .
<i>Copy All Files the First Time the Job Is Run</i>	<p>Initially, this field is set to the value specified in the Default Job settings. You can change it to the value you want to use for this job.</p> <p>If <i>Copy All Files</i> is selected, the job saves versions of all eligible files to the archive database the first time the job is run. Thereafter, the job saves file versions only for files that changed.</p> <p>If <i>Copy All Files</i> is deselected, the job saves versions only of eligible files that were changed during the epoch, such as files that were added, modified, renamed, deleted, or where trustee information changed.</p>
<i>Do Not Run</i>	<p>If <i>Do Not Run</i> is selected (enabled), the job enters the Stopped state when you save the job. You can start or schedule the job at any time thereafter.</p> <p>If <i>Do Not Run</i> is deselected (disabled), the job enters the Scheduled state when you save the job.</p> <p>However, if any setting is missing or invalid, the job enters the Not Configured state when you save the job, regardless of the <i>Do Not Run</i> setting. Edit the job to make corrections.</p>

- 2 When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

7.5.3 Setting Source Server Information

- 1 On the Job Settings page, configure the following source server information properties:

Property	Description
<i>Use Default Job's Source Server</i>	<p>If the default job defines a default source server, the <i>Use Default Job's Source Server</i> check box is selected. The default value is displayed in the field and you cannot modify the value. Deselect <i>Use Default Job's Source Server</i> to modify the values for that field.</p> <p>If the default job does not define a default source server, the <i>Use Default Job's Source Server</i> option is deselected and disabled.</p>
<i>Server</i>	<p>The host name of the source server where the data to be versioned is located. For example, <code>servername.context</code>. The source server is in the same eDirectory tree as the Archive and Version Services server.</p> <p>If you specify a new source server for an existing job, the archived data becomes associated with the source volume on the new source server. Typically, you change the source server only when you have renamed it.</p>
<i>Volume</i>	<p>The name of the source volume (or mount point) where the data to be versioned is located. For example, <code>users</code> or <code>data</code>.</p>
<i>Use Default Job's Snapshot Pool</i>	<p>If the default job defines a default snapshot pool, the <i>Use Default Job's Snapshot Pool</i> check box is selected. The default value is displayed in the field and you cannot modify the value. Deselect <i>Use Default Job's Snapshot Pool</i> to modify the values for that field.</p> <p>If the default job does not define a default snapshot pool, the <i>Use Default Job's Snapshot Pool</i> option is deselected and disabled.</p>
<i>Snapshot Pool</i>	<p>The name of the destination pool where snapshots of the source volume are created and temporarily maintained at the end of an epoch until point-in-time file versions can be saved to the archive database. For example, <code>pusers_sl</code> or <code>pdata_sl</code>.</p> <p>If no snapshot pool name is specified, or if the snapshot cannot be created for any reason, the versioning process copies files directly from the source volume. In this case, open files cannot be versioned.</p>
<i>Free Space ID</i>	<p>ID of the free space object to be used for storing snapshot data. For example, <code>/dev/hda</code>.</p>
<i>Sectors</i>	<p>Specifies the number of sectors on free space to be used for storing snapshot data. The value needs to be an integer, for example, 409600.</p>
<i>Filter</i>	<p>Set filters to specify the types of files and directories to be versioned. For more information, refer to the online help.</p>

- 2 When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

7.5.4 Setting Run Schedule Information

- 1 On the Job Settings page, specify when to start the job and the frequency for running the job. To set the frequency, you must specify only one of three scheduling options:

Property	Description
<i>Use Default Job's Run Schedule</i>	Select <i>Use Defaults</i> to use the <i>Run Schedule</i> settings specified on the Default Job Settings page. If there is no defined default schedule, this option is disabled.
<i>Every</i>	<p>Select <i>Every</i> to enable <i>Scheduled Interval</i>. Specify the elapsed time between the beginning of versioning processes for a job. In the <i>Units</i> drop-down list, select seconds, minutes, hours, or days. For example, 45 seconds, 1 minute, 15 minutes, 2 hours, or 12 hours.</p> <p>If the versioning process exceeds the time specified as the interval, the overlapping scheduled job is skipped. No file versions are saved for skipped job runs. After the version process completes, the job runs at its next scheduled interval. If you observe that the job skips some versioning intervals, you can increase the interval between versions or reduce the amount of data to be versioned.</p>
<i>Time</i>	<p>Select the <i>Time</i> field, then specify the start time when the job's version process begins, then specify one or more days of the week to run the job.</p> <ul style="list-style-type: none"> ♦ Start Time: Click the double-arrows to navigate to the time by hours or click the single-arrow to navigate by 15-minute increments. For example, from 00:00 AM, click the double-arrow to move to 01:00 AM, then click the single arrow three times to reach 01:45 AM. ♦ Days of the Week: Select the <i>Days of the Week</i> check box to select all days for a daily job run, or select only the check boxes next to one or more days of the week you want the job to run. Choices include <i>Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, Sunday, and All</i>.

- 2 When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

7.5.5 Setting Delete Policy Information

The delete policy determines the retention of file versions by age or by number of versions. If a delete policy is set, the job runs a process to delete file versions according to its own Delete Schedule. The delete process is not related to the job's Run Schedule, which determines when file versions are saved from the source volume. The job's delete policy runs if the job is in a Running, Scheduled, or Stopped state. The job's delete policy does not run if a job is in the Clean Up Users, Not Configured, or Deleted state.

IMPORTANT: The Delete Schedule operates separately from the Run Schedule.

- 1 Use the Job Settings page to specify one of the following options:

Property	Description
<i>No Delete Policy</i>	Select this option if you want to retain file versions indefinitely.
<i>Use Default Job's Delete Policy</i>	Select <i>Use Defaults</i> to use the <i>Delete Policy</i> settings specified on the Default Job Settings page. If there is no defined default delete policy, this option is disabled and <i>Keep Latest Version of Current File</i> is selected by default.
<i>Define Job Delete Policy</i>	Specifies the schedule for deleting file versions, if they are eligible for deletion. Set the <i>Interval</i> and <i>Maximum Keep</i> settings.

- 2 If you selected *Define Job Delete Policy*, complete the following information:

Property	Description
<i>Interval</i>	This value represents the amount of time to wait from the time a delete-versions process ends until another delete-versions process begins. If a value is not specified, 24 hours is the default interval. How long the delete process takes depends on the number of files stored in the archive database. For example, suppose you set the Delete Schedule interval to 1 hour. When you save a job, ArkManager starts the interval timer. After 1 hour elapses, the job kicks off its delete-versions process, resets the interval timer, and pauses the timer. When the delete process ends, the interval timer begins. The delete intervals repeat in this manner until the job is stopped.
<i>Maximum Keep</i>	Specifies the maximum number of versions of each file to keep in the archive and how long to keep file versions. At least one of the values must be non-zero. If you set both the <i>Maximum Keep Versions</i> and <i>Maximum Keep Time</i> to zero values, the Delete Policy function does not run. <ul style="list-style-type: none"> ♦ Time: The maximum time that a file version is maintained in the archive. Specify the time in whole numbers; then in the Units drop-down list, select seconds, minutes, hours, or days. ♦ Versions: Specify the maximum number of versions of each file to keep in the archive. As the number of versions exceeds this integer value, the oldest version is deleted. ♦ Keep Current Copy: By default, the latest file version of current files remains in the database, even if the <i>Maximum Keep Time</i> elapses. If <i>Keep Current Copy</i> is selected, the archive keeps the latest file version as long as its source file is current on the source volume, even if the <i>Maximum Keep Time</i> elapses. After the user deletes the current source file, the deletion is noted when the versioning process runs. If the file version's age is within the <i>Maximum Keep Time</i>, the archive database retains the file version of the deleted file; otherwise, the file version is deleted. If <i>Keep Current Copy</i> is deselected, the archive deletes the file version when the <i>Maximum Keep Time</i> elapses.

- 3 When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

7.6 Creating a Job

If your implementation plan calls for multiple archive servers, you must configure each with its own set of archive jobs.

Before you configure individual jobs, make sure you have configured the following information:

- ♦ **Archive Server Information:** Configure basic information for your archive server. For information, see [Section 7.3, “Configuring or Viewing Archive Server Properties,” on page 54](#).
- ♦ **Default Job Settings:** Default settings can optionally be used by any job on the archive server. Each job can optionally use one or more of the defaults that are set; each job’s usage is independent of other jobs.

To use the defaults, you must set default job settings before the job runs. (In *Roles and Tasks*, expand the *Archive Versioning* role, select *Archive Server Properties*, the select *Default Job Settings*.) For information, see [Section 7.4, “Configuring Default Job Settings,” on page 56](#).

To create a job:

- 1 In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
- 2 Select the archive server you want to manage, then wait for the page to refresh.
For information, see [Section 7.2, “Configuring Archive Volume,” on page 54](#).
- 3 Select the *Jobs* tab.
- 4 Click *New* to open the Create New Job page.
- 5 Specify the property settings to use when archiving file versions for a specified volume.

For information about property settings, see [Section 7.5, “Configuring Job Properties,” on page 60](#).

You must at least specify a job property or use the default job property for all required fields before the job can run. Required fields are marked with an asterisk (*).

Roles and Tasks

[All Categories]

Archive Versioning

Archive Jobs

Archive Server Properties

Clusters

DHCP (OES Linux)

Directory Administration

Distributed File Services

DNS

eDirectory Encryption

eDirectory Maintenance

File Access (NetStorage)

File Protocols

Files and Folders

Groups

Help Desk

iFolder 3.7

iPrint

Kerberos Management

LDAP

Linux User Management

NMAAS

Novell Certificate Access

Novell Certificate Server

Partitions and Replicas

Passwords

QuickFinder

Rights

Schema

SMS Backup and Restore

SNMP

Storage

Users

WAN Traffic

New Job

required = *

Define only one job for a volume. Job names must be unique and cannot be reused. Specify the job's property settings, using default values as desired.

Job Information

Name *: Weeklyjob

☒ Copy all files first time the job is run
 ☐ Do not run

Source Server Information

☐ Use Default Job's Source Server

Server *: Archive-server.novell

Mount Point *: ARKVOL

Snapshot

☐ Use Default Job's Snapshot Pool

Snapshot Pool:

Free Space: hda3.freespace1 (FreeSpaceID,Size)

Sectors: 1000000

Note: Number of Sectors should be less than the FreeSpace Size

Filter

☒ Include
 ☐ Exclude

Path:

Extension:

Pattern:

Wildcard:

☐ Override Default Job's Filter

Delete Policy

☐ No Delete Policy
 ☐ Use Default Job's Delete Policy
 ☒ Define Job Delete Policy

Schedule Interval: Days

Maximum Keep: Time: 60 Days

Versions:

☒ Keep latest version of current file

Run Schedule *

☐ Use Default Job's Run Schedule
 ☐ Every: Seconds
 ☒ Time: 22:30 (10:30 PM)

Days

☐ Sunday
 ☐ Monday
 ☐ Tuesday
 ☐ Wednesday
 ☐ Thursday
 ☐ Friday
 ☐ Saturday

OK

Cancel

- 6 When you are done, click *OK* to save your changes, or click *Cancel* at any time to discard them. The job is saved only if all required fields are completed and all settings are valid. If necessary, correct any missing or invalid information, then click *Apply*. For information about annotations for missing or invalid information, see [Section 7.7.1, "Correcting Missing or Invalid Information,"](#) on page 67.
- 7 The *Archive Jobs* window lists the scheduled job.

Jobs					1 Item(s)
New	Edit	Delete	Refresh	Actions	
Name	Status	Next Start Date	Source Server	Volume	
Weeklyjob	Scheduled	5/21/07 3:10 PM	Archive-server-	PRIMEVOL2	

After configuring a job, you can access file versions using NSS File Version Utility or Web-based Archive function using Novell NetStorage.

66 OES 11: Novell Archive and Version Services 2.1 Administration Guide

7.7 Editing a Job's Properties

After you configure a job, you might occasionally need to modify its properties, such as to change the run schedule or delete policy. If you modify the server's default job settings, it can affect jobs that use those defaults.

- ♦ [Section 7.7.1, "Correcting Missing or Invalid Information," on page 67](#)
- ♦ [Section 7.7.2, "Modifying Job Settings," on page 68](#)
- ♦ [Section 7.7.3, "Applying Modified Job Settings," on page 68](#)

7.7.1 Correcting Missing or Invalid Information

For existing jobs, the Job Properties page marks fields that are missing information or contain invalid information.

Required Fields

You must specify a job property or use the default job property for a required field before the job can run. Required fields are marked with an asterisk (*).

Not Configured Fields

The job might report a field as Not Configured if the setting is missing or invalid, or if the default job setting the field uses are missing or invalid. While a job is in the Not Configured state, the job and its delete policy do not run.

Settings that need attention are marked with three asterisks (***) on a red background. Verify the job settings and make any desired changes. If necessary, go to the *Archive Server Properties > Default Job Settings* to edit default job settings.

For example, fields marked as Not Configured might have missing or invalid data under the following circumstances:

- ♦ If you define a job directly in the `arkConfig.xml` file and a required setting is missing or invalid, the job is placed in a Not Configured state. (If you define jobs in the Archive Versioning plug-in to iManager, the interface does not allow you to save a job with invalid or missing values.)
- ♦ If you remove settings in the *Default Job Settings*, jobs that use the defaults are affected.
- ♦ If a source volume is not mounted, it appears that the setting is invalid.
- ♦ If the source server is down, it appears that the setting is invalid.
- ♦ If you previously removed a job's definition (job name) directly in the `arkConfig.xml` file instead of deleting the job with iManager, the job's data (job properties) remains in the database, but it has no job definition.

Defaults

To use the default job properties, you must set default job settings before the job runs. In *Roles and Tasks*, click *Archive Server Properties*, then select *Default Job Settings* to configure default properties.

7.7.2 Modifying Job Settings

- 1 In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
- 2 Select the archive server you want to manage, then wait for the page to refresh.
For information, see [Section 7.2, “Configuring Archive Volume,” on page 54](#).
- 3 Select the *Jobs* tab.
- 4 (Optional) Select *Archive Server Properties*, then select *Default Settings* to verify or set default job settings you want to use for the job.
For information, see [Section 7.4, “Configuring Default Job Settings,” on page 56](#).
- 5 On the *Jobs Report*, select a *Job* check box, then click *Edit* to open the Edit Job page.
- 6 Specify settings for all required fields and for any optional settings you want to use.
For information, see [Section 7.5, “Configuring Job Properties,” on page 60](#).
You must specify a valid job property, or use the default job property, for all required fields before the job can run. Required fields are marked with an asterisk (*).
- 7 Correct any invalid settings.
Invalid information is marked with three asterisks (***) on a red background. For information, see [Section 7.7.1, “Correcting Missing or Invalid Information,” on page 67](#).
- 8 When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.
The job is saved only if all required fields are completed and all settings are valid.

7.7.3 Applying Modified Job Settings

It is not necessary to stop jobs while you modify their settings. If all required and optional settings are valid:

- ♦ The job is rescheduled forward from the save time.
- ♦ If a delete-versions job is in progress, the run is completed with the old delete policy. A modified delete policy is applied when the delete job runs again.
- ♦ You can expect the following additional actions, depending on the Do Not Run setting and the state of the job when you save your changes:

Table 7-2 *How Changes to Job Settings Take Effect*

Do Not Run Setting	State of the Job When You Save Edits	Additional Actions
Enabled (selected)	Running	The running job is completed with the old settings, then the job enters the Stopped state.
	Scheduled, Stopped, or Not Configured	The job enters the Stopped state.
	Clean Up Users	An in-progress cleanup of the job's users stops gracefully when you begin the edit, and the job enters the Stopped state. When you save changes, the job remains in the Stopped state.
Disabled (deselected)	Running	The running job is completed with the old settings, then it enters the Scheduled state.
	Scheduled, Stopped, or Not Configured	The job enters the Scheduled state.
	Clean Up Users	An in-progress cleanup of the job's users stops gracefully when you begin the edit, and the job enters the Stopped state. When you save changes, the job enters the Scheduled state.

7.8 What's Next

To manage Archive and Version services jobs, see [“Managing Jobs” on page 71](#).

8 Managing Jobs

After creating jobs for Novell Archive and Version Services 2.1, manage them with the Archive Versioning plug-in for Novell iManager.

This section discusses the following tasks:

- ♦ [Section 8.1, “Viewing a Jobs Report,” on page 71](#)
- ♦ [Section 8.2, “Viewing a Job’s Details,” on page 73](#)
- ♦ [Section 8.3, “Starting or Scheduling a Job,” on page 74](#)
- ♦ [Section 8.4, “Stopping a Job,” on page 75](#)
- ♦ [Section 8.5, “Viewing the Archive Log,” on page 76](#)
- ♦ [Section 8.6, “Filtering Messages in the Archive Log,” on page 78](#)
- ♦ [Section 8.7, “Cleaning Up the Job’s User List,” on page 79](#)
- ♦ [Section 8.8, “Deleting File Versions,” on page 80](#)
- ♦ [Section 8.9, “Deleting a Job,” on page 81](#)
- ♦ [Section 8.10, “Viewing a Deleted Jobs Report,” on page 81](#)
- ♦ [Section 8.11, “Salvaging a Deleted Job,” on page 82](#)
- ♦ [Section 8.12, “Purging a Deleted Job,” on page 83](#)

8.1 Viewing a Jobs Report

The jobs report lists all current jobs on a selected archive server and reports their status and schedule.

- ♦ [Section 8.1.1, “Generating the Jobs Report,” on page 71](#)
- ♦ [Section 8.1.2, “Sorting the Jobs Report by Column,” on page 72](#)
- ♦ [Section 8.1.3, “Setting the Refresh Rate for the Jobs Report,” on page 72](#)
- ♦ [Section 8.1.4, “Understanding the Report Content,” on page 72](#)

8.1.1 Generating the Jobs Report

- 1 In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
- 2 Select the archive server you want to manage, then wait for the page to refresh.
For information, see [Section 7.2, “Configuring Archive Volume,” on page 54](#).
- 3 Select the *Jobs* tab.
- 4 View a report of all its current Archive and Version Services jobs.
Each entry in the table represents a single job.

8.1.2 Sorting the Jobs Report by Column

All columns are sortable in ascending and descending order. Click a heading link to sort the jobs by that column. Click the link a second time to sort the jobs in reverse order. A sort icon next to the heading indicates which column is being used as the sort key and the sort order.

8.1.3 Setting the Refresh Rate for the Jobs Report

You can control the refresh frequency of the jobs report. The setting persists only while you are viewing the report or while the browser session is current.

Select *Refresh*, then select the desired frequency from the drop-down list:

Off
Immediately
Every 5 seconds
Every 15 seconds
Every 30 seconds [Default]
Every 60 seconds
Every 5 minutes
Every 15 minutes

8.1.4 Understanding the Report Content

The following table describes information in the Jobs Report:

Table 8-1 Description of Information in the Jobs Report

Property	Description
<i>Jobs</i>	Select a <i>Job</i> check box to select the job you want to manage.
<i>Name</i>	The administrator-specified name of the job. Typical names might provide information about the source server and source volume. For example: <i>server_volume</i> or <i>srv1_vol1</i> . Click the job's <i>Name</i> link to edit the job.
<i>Status</i>	Reports the current operational state of the job. Click the <i>Status</i> icon to view the job's details. <ul style="list-style-type: none">• Scheduled: The job is scheduled to save file versions; it is not currently running.• Running: The job is actively saving file versions to the archive database.• Stopped: The job's archived data is available to users, but the job is disabled so that no new versions are being captured.• Cleanup: The Clean Up Users job is in progress. For information, see Section 8.7, "Cleaning Up the Job's User List," on page 79.• Not Configured: One or more of the job's settings are missing or invalid, or the default job settings the job uses are missing or invalid. While a job is in the Not Configured state, the job and its Delete policy do not run.

Property	Description
<i>Status</i> (cont'd)	<p>Settings that need attention are marked with three asterisks (***) on a red background. Click the job's <i>Name</i> link to view or edit the job properties. Go to the Default Job Settings page (Archive Versioning > Archive Server Properties > Default Job Settings) to edit default job settings. Verify the job settings and make any necessary changes.</p> <p>For example, fields marked as Not Configured might have missing or invalid data under the following circumstances:</p> <ul style="list-style-type: none"> ♦ If you define jobs in the Archive Versioning plug-in to iManager, the interface does not allow you to save a job with invalid or missing values. ♦ If you remove settings in the Default Job Settings, jobs that use the defaults are affected. ♦ If a source volume is not mounted, it appears that the setting is invalid. ♦ If the source server is down, it appears that the setting is invalid. ♦ If you previously removed a job's definition directly in the <code>/etc/opt/novell/arkManager/conf/arkConfig.xml</code> file instead of deleting the job with iManager, the job's data remains in the database, but it has no job definition.
<i>Next Start Date</i>	<p>The next scheduled date and time that the job is scheduled to begin a versioning process. The format is mm/dd/yy hh:mm xm. For example: 4/25/05 11:45 PM.</p> <p>If the job is in a Stopped state, the <i>Next Start Date</i> might have a date and time that is previous to the current time. Whenever the job starts again, you can specify whether to run the job immediately (<i>Run Now</i> or <i>Run Now with Copy All</i> actions) or to schedule the job to run later (<i>Run Later</i> or <i>Run Later with Copy All</i> actions).</p>
<i>Source Server - Source Volume</i>	<p>The server name and volume name of the volume that the job archives. A source volume can have only one archive job defined for it.</p>

8.2 Viewing a Job's Details

The Job Details page shows the effective settings for a job, which are a combination of the job's individual properties and the default settings the job uses.

- 1 In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
- 2 Select the archive server you want to manage, then wait for the page to refresh.
For information, see [Section 7.2, "Configuring Archive Volume," on page 54](#).
- 3 Select the *Jobs* tab.
- 4 View a report of all its current Archive and Version Services jobs.
Each entry in the table represents a single job.
- 5 Use one of the following methods to view the details of an Archive Job:
 - ♦ Select a *Job* check box, then click *Actions > Details*. Select only one job at a time.
 - ♦ Click the *Status* icon for the job.

The *Job Details* box displays the following information about the selected job.

Property	Description
<i>Name</i>	The administrator-specified name of the job.
<i>Status</i>	<p>Reports the current operational state of the job.</p> <ul style="list-style-type: none"> ■ <i>Scheduled</i>: The job is scheduled to save file versions; it is not currently running. ■ <i>Running</i>: The job is actively saving file versions to the archive database. ■ <i>Stopped</i>: The job's archived data is available to users, but the job is disabled so that no new versions are being captured. ■ <i>Cleanup</i>: The Clean Up Users job is in progress. For information, see Cleaning Up a Job's Users List. ■ <i>Not Configured</i>: One or more of the job's settings are missing or invalid, or the default job settings the job uses are missing or invalid. While a job is in the <i>Not Configured</i> state, the job and its Delete policy do not run.
<i>Source Server</i>	The eDirectory Server object name of the source server where the data to be versioned is located. The source server is in the same tree as the Archive and Version Services server.
<i>Source Volume</i>	The name of the volume where the data to be versioned for this job is located.
<i>Snapshot Pool</i>	The name of the destination pool where snapshots of the source volume are created and temporarily maintained until point-in-time file versions can be saved to the archive database.
<i>Last Start Date</i>	The date and time that the job started the last time it was successfully run and completed. If the job is in the Running state, this time stamp is not updated until the job completes successfully.
<i>Next Start Date</i>	<p>The next scheduled date and time that the job is scheduled to begin a versioning process. The format is mm/dd/yy hh:mm xm. For example: 04/25/05 11:45 PM.</p> <p>If the job is in a Stopped state, the <i>Next Start Date</i> might have a date and time that is previous to the current time. Whenever the job starts again, you can specify whether to run the job immediately (<i>Run Now</i> or <i>Run Now with Copy All</i> actions) or to schedule the job to run later (<i>Run Later</i> or <i>Run Later with Copy All</i> actions).</p>
<i>Schedule</i>	<p>Reports the interval or start times for the selected job.</p> <ul style="list-style-type: none"> ◆ Scheduled Interval: The elapsed time between the beginning of successive versioning processes for a job. ◆ Scheduled Start Time: The time of day when the job's version process begins and the days of the week the job runs.

6 When you are done, click *Close*.

8.3 Starting or Scheduling a Job

- 1 In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
- 2 Select the archive server you want to manage, then wait for the page to refresh.
For information, see [Section 7.2, "Configuring Archive Volume," on page 54](#).
- 3 Select the *Jobs* tab.
- 4 View a report of all its current Archive and Version Services jobs.

Each entry in the table represents a single job.

- 5 In the *Jobs Report*, select one or more *Job* check boxes, click *Actions*, then select one of the following actions:

Property	Description
<i>Run Now</i>	Starts the selected jobs immediately.
<i>Run Later</i>	Schedules the selected jobs to run at their individual next scheduled start times.
<i>Run Now with Copy All</i>	Starts the selected jobs immediately and copies all files specified by the individual selected jobs from their source volumes to the archive volume on a one-time, special version occurrence. NOTE: This option takes precedence over the individual Job properties set for copying the files the first time.
<i>Run Later with Copy All</i>	Schedules the selected jobs to run at their individual next scheduled start times. For that run only, it copies all files specified by the individual selected jobs from their source volumes to the archive volume on a one-time, special version occurrence.

8.4 Stopping a Job

Stop a job whenever you want to temporarily pause versioning or to cease versioning a particular source volume.

- 1 In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
- 2 Select the archive server you want to manage, then wait for the page to refresh.
For information, see [Section 7.2, “Configuring Archive Volume,” on page 54](#).
- 3 Select the *Jobs* tab.
- 4 View a report of all its current Archive and Version Services jobs.
Each entry in the table represents a single job.
- 5 In the *Jobs Report*, select one or more *Job* check boxes, then click *Actions > Stop*.

Stopping a job stops the selected jobs gracefully and disables them from running again until you start or schedule each job. The Stopped setting persists through a server reboot. For information, see [Section 8.3, “Starting or Scheduling a Job,” on page 74](#).

Stopping a job does not stop the delete policy from running. To stop the delete policy, modify the job’s *Delete Policy* setting to *No Delete Policy*. For information, see [Section 7.7, “Editing a Job’s Properties,” on page 67](#).

You can also use the `<stopped/>` tag in the *Job* element for the particular job in the `arkConfig.xml` file.

8.5 Viewing the Archive Log

ArkManager issues messages as it performs tasks for each job. The Archive Log displays error, warning, and normal messages reported by all jobs on a selected archive server. You can filter log entries to view only those messages of interest (such as by severity type or by job), to set the number of messages to view per page, and to specify the date of the messages you want to jump to first in the log.

This section discusses the following topics:

- [Understanding Log Entries](#)
- [Viewing Messages in the Archive Log](#)
- [Viewing Messages in the Server Log File](#)

8.5.1 Understanding Log Entries

The Archive Log contains all messages for selected jobs on a selected archive server. By default, the log displays messages for all severity types and for all jobs, beginning with the current date and time. Each log entry displays the following information:

Table 8-2 Description of Parameters in the Archive Log

Property	Description
<i>Job Name</i>	The administrator-specified unique name of the job.
<i>Severity</i>	Specifies the alert level of the job's message. Messages have three levels of severity: <ul style="list-style-type: none">♦ Error: Reports a failure to perform a job task. For example, the versioning service might not be able to save a version of a file from the source volume to the archive volume because of a broken connection.♦ Warning: Reports a non-critical error in performing a job.♦ Normal: Reports the status of normal job tasks such as starting, copying, and completing.
<i>Date</i>	The date and time stamp of the message. The format is mm/dd/yy hh:mm xm. For example: 04/25/05 11:45 PM.
<i>Message</i>	The text message issued by the job.

8.5.2 Viewing Messages in the Archive Log

- [“Viewing a Log of All Jobs” on page 77](#)
- [“Viewing a Log of Selected Jobs” on page 77](#)
- [“Navigating the Archive Log” on page 77](#)
- [“Updating the Log Message Entries” on page 77](#)
- [“Filtering the Log Message Entries” on page 77](#)

Viewing a Log of All Jobs

- 1 In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
- 2 Select the archive server you want to manage, then wait for the page to refresh.
For information, see [Section 7.2, “Configuring Archive Volume,” on page 54](#).
- 3 Select the *Log* tab.
- 4 View the Archive and Version Services log, which contains all messages for all jobs on the selected server. You can apply a filter for message types or time periods.

Viewing a Log of Selected Jobs

- 1 In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
- 2 Select the archive server you want to manage, then wait for the page to refresh.
For information, see [Section 7.2, “Configuring Archive Volume,” on page 54](#).
- 3 Select the *Jobs* tab.
- 4 In a Jobs report, select one or more *Job* check boxes, then click *Actions > View Log*.
- 5 View the Archive and Version Services log.
The log opens with a filter applied so that the log displays only messages for the selected jobs.

Navigating the Archive Log

By default, the log displays 10 messages per screen. Messages appear from newest (top and left) to oldest (bottom and right). Use the browser scroll bar to view all entries in the page. Use the arrows at the bottom of the page to move from page to page.

- ♦ Use the left-arrow to navigate to the previous page of newer messages.
- ♦ Use the double left-arrow to go directly to the newest messages.
- ♦ Use the right-arrow to navigate to the next page of older messages.
- ♦ Use the double right-arrow to go directly to the oldest messages.

Updating the Log Message Entries

To view the updated log, click *Archive Job*, then click the *Log* tab. If you have already set the *Refresh Rate* the log file will get updated automatically.

Filtering the Log Message Entries

- ♦ **Filter:** You can view all available log entries, or use a filter to view only those messages you want to see. The log filter allows you to filter messages by severity types and by one or more jobs so that only the messages you want to view are displayed.

By default, the log displays 10 messages per page and begins with the newest messages. You can use the filter to set preferences for the number of log entries to display per page and the date and time of the entries you want to see first in the log.

To set the filter, click *Filter*. For information, see [Section 8.6, “Filtering Messages in the Archive Log,” on page 78](#).

- ♦ **Clear Filter:** Resets the current filter settings to the default log display values, using the current date and time as the starting point of the log, and it displays 10 entries per page.

To reset the filter to the default *Archive Log* values, click *Clear Filter*.

8.5.3 Viewing Messages in the Server Log File

In addition to viewing the log in iManager, you can view logs in the log file at `/var/opt/novell/arkmanager/logs` folder.

8.6 Filtering Messages in the Archive Log

In the iManager Archive Log page, click *Filter* to filter messages so that the log displays only those log entries that you want to see. You can also set preferences for how many log entries to view per page and for the date and time of messages you want to jump to first. You can set one or more of the filter types each time you activate the filter process. Click *Cancel* at any time to back out of the filtering process.

The Filter page allows the following:

- ♦ [Section 8.6.1, “Filtering Log Entries by Severity Type,” on page 78](#)
- ♦ [Section 8.6.2, “Filtering Log Entries by Job,” on page 78](#)
- ♦ [Section 8.6.3, “Setting the Display Details,” on page 79](#)
- ♦ [Section 8.6.4, “Setting the Date and Time Range,” on page 79](#)
- ♦ [Section 8.6.5, “Resetting the Log Filters,” on page 79](#)

8.6.1 Filtering Log Entries by Severity Type

Set filters so that the log displays only the messages that match one or more *Severity Types* that you select. Deselected types are not displayed.

- 1 On the Archive Log page, click *Filter*.
- 2 On the Filter Archive Log page, select one or more *Severity Type* check boxes next to the message types you want to display in the log.
- 3 (Optional) Set any other filter options on the page.
- 4 Click *OK* to save the settings, or click *Cancel* at any time to back out of the filter setting process.

8.6.2 Filtering Log Entries by Job

The archive log displays recent messages for all jobs on the selected server. A filtered log displays messages only for the jobs you select. Generic messages not allocated to any particular job will be displayed under the job name *ArkManager job*.

- 1 On the Archive Log page, click *Filter*.
- 2 Select one or more *Job* check boxes next to the jobs you want to display in the log.
- 3 (Optional) Set any other filter options on the page.
- 4 Click *OK* to save settings, or click *Cancel* at any time to back out of the filter setting process.

8.6.3 Setting the Display Details

- 1 On the Archive Log page, click *Filter*.
- 2 Specify the number of log entries to display per page.
The default is 10 messages.
- 3 Specify one of the following:
 - ♦ Select *Display Log Messages* to display message text in the *Messages* column.
 - ♦ Deselect *Display Log Messages* if you want to omit the *Messages* column in the archive log.
- 4 (Optional) Set any other filter options on the page.
- 5 Click *OK* to save settings, or click *Cancel* at any time to back out of the filter setting process.

8.6.4 Setting the Date and Time Range

The log automatically displays all job messages with the newest messages first. You can specify the date and time of the first message to display. The log jumps directly to the log entries from that date and time instead of to the newest messages. All messages are still available by using the arrow keys at the bottom of the screen to navigate to newer and older messages than the date you set.

To view the most recent log entries:

- 1 Click *View Newest Log Entries*.

To specify a particular date and time for entries to jump to first:

- 1 On the Archive Log page, click *Filter*.
- 2 To specify a date, click the *Calendar* icon, then select a date by using the arrows to navigate to the correct month and year.
- 3 To specify a time, click the arrows to navigate by hours (double arrows) or by 15-minute increments (single arrows) to the time.
- 4 (Optional) Set any other filter options on the page.
- 5 Click *OK* to save settings, or click *Cancel* at any time to back out of the filter setting process.

8.6.5 Resetting the Log Filters

The filter settings are not persistent. They are active only until you use one of the following methods to reset them:

- ♦ Use the *Filter* to specify new *Filter* settings.
- ♦ Click *Clear Filter* on the Archive Log page to return to default settings.
- ♦ Log out of iManager.
- ♦ Close the iManager browser window.

8.7 Cleaning Up the Job's User List

For each job, ArkManager maintains information about its users centrally instead of storing duplicate information among the file versions. The user list grows cumulatively; it is not automatically synchronized to remove users from the list who no longer have associations with current file versions

in the archive. To clear out obsolete user information, run the Clean Up Users task about once a month, depending on the job's schedule, the job's Delete policy, and the turnover of users in your workplace.

During the cleanup, jobs are in the Clean Up Users state. While Clean Up Users is running, the job itself does not run and the job's Delete policy does not run. Afterwards, the jobs are automatically placed in the Stopped state and the jobs' Delete policies are resumed.

If a cleanup job is interrupted, the job stops gracefully and enters the Stopped state. The user list might be only partially cleaned up. For example, the job is interrupted if you delete, edit, start, or schedule the job, or if you shut down or restart ArkManager. If you are editing a cleanup job, when you save your edits, the job continues in the Stopped state if you enabled the *Do Not Run* setting, or it enters the Scheduled state if you disabled the *Do Not Run* setting.

To clean up the user list:

- 1 In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
- 2 Select the archive server you want to manage, then wait for the page to refresh.
For information, see [Section 7.2, "Configuring Archive Volume," on page 54](#).
- 3 Select the *Jobs* tab.
- 4 In the Jobs report, select one or more jobs where you want to clean up the user list, then select *Actions > Stop*.
Wait for the jobs to stop gracefully and report a Stopped status.
- 5 Select one or more jobs, then select *Actions > Clean Up Users*.
When the cleanup for a job is complete, the job enters the Stopped state.
- 6 When cleanup is done, restart the stopped jobs.
For information, see [Section 8.3, "Starting or Scheduling a Job," on page 74](#).

8.8 Deleting File Versions

You can select the file and delete all the versions from the archive database. deleting the file provides free space on the archive volume.

Administrators can delete file versions for a variety of reasons, such as the following:

- ♦ To clear the database of unnecessary data that was versioned prior to applying job filters
- ♦ To delete file versions without compromising the ArkSQL data structure
- ♦ To clear the database of file versions related to a job you plan to delete (See also [Section 8.9, "Deleting a Job," on page 81](#).)
- ♦ To clear the database of older file versions

Deleting a Specified File or Directory Version Set

Use the NSS File Versions Utility or the NetStorage Archive function to locate and delete a selected file version set or a selected directory version and its contents from the archive database. For information, see [OES 11: Novell Archive and Version Services 2.1 User Guide](#).

Deleting Older File Versions

Modify the *Maximum Keep Time*, *Maximum Keep Versions*, or *Keep Current Copy* settings in the job's *Delete Policy* settings to delete only the older file versions in the archive database. For information, see [Section 7.7, "Editing a Job's Properties," on page 67](#).

Deleting All File Versions for a Job

Delete and purge the job to remove all of its file versions from the archive database. For information, see the following:

- ♦ [Section 8.9, "Deleting a Job," on page 81](#)
- ♦ [Section 8.11, "Salvaging a Deleted Job," on page 82](#)

8.9 Deleting a Job

- 1 In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
- 2 Select the archive server you want to manage, then wait for the page to refresh.
For information, see [Section 7.2, "Configuring Archive Volume," on page 54](#).
- 3 Select the *Jobs* tab.
- 4 In the Jobs Report, select one or more *Job* check boxes, click *Delete*, then confirm the deletion.
If the job is running when it is deleted, ArkManager stops the job gracefully, then deletes the job. A partial data set exists for that run, up to the point where the job was stopped.

The deleted job is retained in the archive database in the Deleted job state, but it no longer appears in the list of current jobs in the Jobs Report.

While a job is in the Deleted state:

- ♦ The deleted job does not run.
- ♦ Users cannot access the deleted job's file versions in the archive database.
- ♦ The deleted job's Delete Policy is not enforced.

A deleted job's properties and archived file versions remain in the archive database until you purge or salvage the job. For information, see the following:

- ♦ [Section 8.10, "Viewing a Deleted Jobs Report," on page 81](#)
- ♦ [Section 8.11, "Salvaging a Deleted Job," on page 82](#)
- ♦ [Section 8.12, "Purging a Deleted Job," on page 83](#)

8.10 Viewing a Deleted Jobs Report

A deleted job's properties and archived file versions remain in the archive database until you purge the job. While a job is in the Deleted state:

- ♦ The deleted job does not run.
- ♦ Users cannot access the deleted job's file versions in the archive database.
- ♦ The deleted job's Delete Policy is not enforced.

To view a list of deleted jobs:

- 1 In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
- 2 Select the archive server you want to manage, then wait for the page to refresh.
For information, see [Section 7.2, “Configuring Archive Volume,” on page 54](#).
- 3 Select the *Deleted Jobs* tab.
- 4 View a report of all deleted jobs on the selected server.
- 5 Each entry in the table represents a single job. If multiple deleted jobs have the same job name, they are distinguishable by their deleted times, source servers, and source volumes.

All columns are sortable in ascending and descending order. Click a heading link to sort the jobs by that column. Click the link a second time to sort the jobs in reverse order. A sort icon next to the heading indicates which column is being used as the sort key and the sort order.

Property	Description
<i>Jobs</i>	Select a <i>Job</i> check box to select the job you want to manage.
<i>Name</i>	The administrator-specified name of the job. Typical names might provide information about the source server and source volume. For example: server_volume or srv1_vol1. Click the job's <i>Name</i> link to edit the job.
<i>Source Server - Source Volume</i>	The server name and volume name of the volume that the job archived.
<i>Date</i>	The date and time when the deletion process was completed successfully.

- 6 When you are done, click *Close*.

8.11 Salvaging a Deleted Job

You can salvage a deleted job to restore the job and to allow users to access its archived file versions. You can salvage only one job at a time. You cannot salvage a deleted job if the source volume it archived is currently the source volume for an existing job on the archive versioning server.

- 1 In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
- 2 Select the archive server you want to manage, then wait for the page to refresh.
For information, see [Section 7.2, “Configuring Archive Volume,” on page 54](#).
- 3 Select the *Deleted Jobs* tab.
- 4 In the Deleted Jobs Report, select the *Job* check box next to the job you want to salvage, then click *Salvage*.
If multiple deleted jobs have the same job name, they are distinguishable by their deleted dates, source volumes, and source servers.
- 5 Specify a name for the salvaged job that is unique among the non-deleted jobs on the archive server.
- 6 Click OK.

Salvaging a deleted job restores the job. The following table describes the state of the job after it is salvaged, given the job's state when it was deleted.

State of the Job When Deleted	Delete Action	State of the Job After Salvage
Running	The job stops gracefully, enters the Stopped state, then it is deleted. A partial data set exists for the job.	Scheduled The job's next run time is calculated forward from the time the job is salvaged.
	The job's Delete policy is suspended.	The job's Delete policy runs immediately.
Stopped	The job is deleted.	Stopped
	The job's Delete policy is suspended.	The job remains in the Stopped state until you start or schedule it. The job's Delete policy runs immediately.
Scheduled	The job is deleted.	Scheduled
	The job's Delete policy is suspended.	The job's next run time is calculated forward from the time the job is salvaged. The job's Delete policy runs immediately.
Not Configured	The job is deleted.	Not Configured
	The job's Delete policy, which does not run in the Not Configured state, continues to be suspended.	The job cannot be started or scheduled until you provide valid settings. The job's Delete policy does not run until you start or schedule the job.
Clean Up Users	The Clean Up Users job stops gracefully, enters the Stopped state, then it is deleted. The user list might be only partially cleaned up.	Stopped
	The job's Delete policy, which does not run in the Clean Up Users state, is now suspended.	The job remains in the Stopped state until you start or schedule it. The job's Delete policy runs immediately.

8.12 Purging a Deleted Job

Purging a deleted job removes the job properties and all archived data for the job from the archive database. Purged jobs cannot be recovered.

- 1 In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
- 2 Select the archive server you want to manage, then wait for the page to refresh.
For information, see [Section 7.2, "Configuring Archive Volume," on page 54](#).
- 3 Select the *Deleted Jobs* tab.
- 4 In the Deleted Jobs Report, select the *Job* check box next to the job you want to purge, then click *Purge*.

If multiple deleted jobs have the same job name, they are distinguishable by their deleted dates, source volumes, and source servers.

- 5** Click *OK* to confirm the purge, or click *Cancel* to back out of the process.

9 Running Archive and Version Services in a Virtualized Environment

Archive and Version Services 2.1 runs in a virtualized environment just as it does on a physical NetWare server, or on a physical server running OES 11, and requires no special configuration or other changes.

To get started with virtualization, see [Introduction to Xen Virtualization \(http://www.suse.com/documentation/sles11/book_xen/?page=/documentation/sles11/book_xen/data/cha_xen_basics.html\)](http://www.suse.com/documentation/sles11/book_xen/?page=/documentation/sles11/book_xen/data/cha_xen_basics.html) in the [Virtualization with Xen \(http://www.suse.com/documentation/sles11/book_xen/?page=/documentation/sles11/book_xen/data/book_xen.html\)](http://www.suse.com/documentation/sles11/book_xen/?page=/documentation/sles11/book_xen/data/book_xen.html) guide.

For information on setting up virtualized OES 11, see “[Installing, Upgrading, or Updating OES on a VM](#)” in the *OES 11: Installation Guide*.

10 Configuring Archive and Version Service for Novell Cluster Services

This section describes how to install and configure Archive and Version Services in an cluster environment.

- ♦ [Section 10.1, “Benefits of Configuring Archive and Version Service for High Availability,” on page 87](#)
- ♦ [Section 10.2, “Prerequisites,” on page 87](#)
- ♦ [Section 10.3, “Installing and Configuring Archive and Version Services,” on page 88](#)
- ♦ [Section 10.4, “Novell Cluster Services Configuration and Setup,” on page 90](#)
- ♦ [Section 10.5, “Configuring Jobs using iManager,” on page 95](#)
- ♦ [Section 10.6, “Managing Archive Resource,” on page 95](#)

10.1 Benefits of Configuring Archive and Version Service for High Availability

Novell Archive and Version Services provides a convenient and cost-effective way for individual users to instantly restore previous versions of modified, renamed, or deleted network files. If for some reason the Archive and Version service is not accessible, clients lose their ability to retrieve archived versions of their files. The Archive service is automatically started, stopped, and restarted on different servers in the cluster by Novell Cluster Services. Hence, running Archive and Version Services in a cluster environment is beneficial as Archive service is highly available and you can retrieve the versions of files even if the Archive server is unavailable on a particular node.

Before you attempt to implement this solution, familiarize yourself with how Cluster Services works. For information, see the [OES 11: Novell Cluster Services 2.0 for Linux Administration Guide](#).

10.2 Prerequisites

- ♦ eDirectory
- ♦ Novell Cluster Services (NCS)
- ♦ Novell Storage Services (NSS)

10.3 Installing and Configuring Archive and Version Services

During OES installation, select Archive and Version Services and Novell Cluster Services options. For Archive and Version Services, select the *Configure Later* option. You must install Novell Cluster Services (NCS) on the OES 11 server before configuring Archive and Version Services to work in a cluster environment. NCS provides a Archive and Version Services (AV) resource template, which facilitates configuring the service in a cluster environment.

For highly availability of Archive service, you need to enable cluster service on Archive volume and PostgreSQL database. The Archive volume needs to be on an NSS file system and PostgreSQL database directories on Linux POSIX Volumes.

After you have installed, set up, and configured Novell Cluster Services, you need to create [NSS](#) and [Linux POSIX](#) shared volumes before configuring Archive and Version service to run in a cluster environment.

10.3.1 Creating NSS Shared Disk Partition

To make the Archive volume highly available, you must create a shared NSS pool and volume on a shared device that can be migrated or failed over to other nodes in a cluster. The shared volume is the location you want to use for the Archive volume.

- ♦ [“Primary Volume \(Optional\)” on page 88](#)
- ♦ [“Archive Volume” on page 88](#)
- ♦ [“Making the Volume Higly Available” on page 89](#)

Primary Volume (Optional)

You must create a shared disk partition for an Primary Volume, if you need to make the Primary Volume highly available.

Create a NSS shared disk partition for an Primary Volume and cluster enable it. For more information, see [“Configuring Cluster Resources for Shared NSS Pools and Volumes”](#) in the *OES 11: Novell Cluster Services 2.0 for Linux Administration Guide*.

Archive Volume

You must create a shared disk partition for an Archive Volume and cluster enable it. For more information, see [“Configuring and Managing Cluster Resources for LVM Volume Groups”](#) in the *“OES 11: Novell Cluster Services 2.0 for Linux Administration Guide”*.

NOTE: Cluster resource [load](#) and unload scripts are automatically generated for pools when they are cluster-enabled.

Making the Volume Highly Available

Using iManager, proceed with the following steps to make the Archive volume highly available.

This procedure uses the following sample values for the virtual cluster server name of the cluster resource you created for the NSS file system, the server context, a LUM group for the users who want to restore file versions, and a username. Make sure to change these values to the actual values in your solution.

- ♦ Virtual Cluster Server Name: Archive-server
 - ♦ Server Context: finance
 - ♦ Username: Carlos
 - ♦ LUM group: Archivegp
- 1 Click the *Linux User Management > Create Unix Workstation Object* task. This opens the Create Unix Workstation Object window.
 - 1a In the *Unix Workstation Name* field, specify Archive-server.
 - 1b In the *Context for Unix Workstation Object* field, search and select *finance*.
 - 1c In the *Unix Config Object* field, search and select *UNIX Config. finance*.

- 2 Click Next to create a Unix workstation.

A *UNIX WorkStation-Archive-server.finance* workstation is created for Primary Volume.

- 3 The user who needs to restore the file versions needs to be LUM-enabled. To do this, click the *Linux User Management > Enable Users for Linux* task.

- 3a In the *Select Users* window, select *Carlos* and click *Next*.
- 3b Every Linux user must belong to a primary group. In the *Select Primary Group* window, select *novlxtier* object and click *Next*.

IMPORTANT: A user must be a part of the *novlxtier* group to use the NSS File Version Utility for restoring versions.

- 3c In the *Select Workstations* window, select *UNIX WorkStation-Archive-server.finance* and click *Next*.

The user *Carlos* is LUM-enabled and his primary group is *novlxtier*. *Carlos* is also given access to *UNIX WorkStation-Archive-server.finance* workstation.

NOTE: In addition to *novlxtier* group, user can be added to another groups using iManager. To add a user to *Archivegp* group, click *Users>Modify User* task. Select the user *Carlos* and click *OK*. Select the tab *Group Memberships* and add *Archivegp* group.

The user *Carlos* is now a member of *novlxtier* and *Archivegp* groups.

10.3.2 Creating Linux POSIX File System Shared Disk Partitions

The PostgreSQL database's data directories need to be on Linux POSIX file system volumes. Create and mount the Linux POSIX Volumes. For more information, see [“Configuring and Managing Cluster Resources for LVM Volume Groups”](#) in the *OES 11: Novell Cluster Services 2.0 for Linux Administration Guide*.

On mounting the Linux POSIX Volumes, continue with [Section 10.3.3, “Configuring Archive and Version Service on the Shared Disk,”](#) on page 90.

NOTE: You must cluster enable the Linux POSIX Volumes, see [Section 10.4.2, “Creating an Archive Service Cluster Resource,”](#) on page 91 for more information.

10.3.3 Configuring Archive and Version Service on the Shared Disk

To configure Archive service to use shared storage, you need to create a shared directory (file system or disk) on the shared disk system and create mount points to that shared file system on each cluster server that will run Archive service.

- 1 Edit the `arkdatadir.conf` file located in `/etc/opt/novell/arkmanager/conf/` and change the existing path from `/var/opt/novell/arkmanager/data` to `MOUNT_POINT/novell/arkmanager/data`.

Where `DATA_PATH` points to location of Linux POSIX volumes on a shared disk and `MOUNT_POINT` points to mounted path of the Linux POSIX volumes on a shared disk. The value of `MOUNT_POINT` is same as used in the Archive load script. For example, `MOUNT_POINT = /mnt/ark`

- 2 At the command prompt, execute the following command:

```
yast2 arkmanager
```

Authenticate with admin password. Enter the inputs for the fields database port, username and password, refer help on the left pane of the screen for details on the fields.

To configure Archive and Version Service, refer [Section 5.2, “Installing and Configuring an Archive Server,”](#) on page 41.

- 3 To configure Archive and Version services on each cluster node, repeat [Step 1](#) and [Step 2](#).

10.4 Novell Cluster Services Configuration and Setup

After Archive service is properly installed and configured, you must create and configure Archive and Version service cluster resource in Novell Cluster Services. This includes configuring Archive and Version service load and unload scripts; setting Archive and Version service start, failover, and failback modes; and assigning the Archive and Version service resource to specific servers in your cluster.

10.4.1 Archive and Version Resource Template

Novell Cluster Services includes an Archive and Version service resource template (`AV_Template`), that greatly simplifies the process for creating an Archive and Version service cluster resource. Much of the Archive and Version service cluster resource configuration is performed automatically by the `AV_Template`.

- 1 Open your Internet browser and enter the URL for iManager.
The URL is `https://server_ip_address/nps/iManager.html`. Replace `server_ip_address` with the IP address or DNS name of a server in the cluster or with the IP address for Apache-based services.
- 2 Enter your username and password.
- 3 In the left column, locate and click *Clusters > Cluster Options*. You can configure and manage cluster using this section.
- 4 Specify a cluster name or browse and select one, the page displays cluster resources for the selected cluster.

5 The *Cluster Objects*, lists *AV_Template*.

5a Click *AV_template > Scripts*

The Load Script includes commands to start the Archive service, the Unload Script to stop the Archive server, and the Monitor script allows Novell Cluster Services to detect when an individual resource on a node has failed independent of its ability to detect node failures.

10.4.2 Creating an Archive Service Cluster Resource

Novell Cluster Services includes an Archive and Version service resource template (*AV_Template*), that greatly simplifies the process for creating an Archive and Version service cluster resource. Much of the Archive and Version service cluster resource configuration is performed automatically by the *AV_Template*.

For more information on creating a cluster resource, see “[Configuring and Managing Cluster Resources for LVM Volume Groups](#)” in the *OES 11: Novell Cluster Services 2.0 for Linux Administration Guide*.

NOTE: Cluster resource load and unload scripts are automatically generated for NSS pools when they are cluster-enabled.

To create an Archive and Version service cluster resource:

- 1 Offline the cluster resource to stop the service and unmount the shared disk (file system) that was created in [Section 10.3.2, “Creating Linux POSIX File System Shared Disk Partitions,” on page 89](#).

For information, see “[Onlining and Offlining \(Loading and Unloading\) Cluster Resources from a Cluster Node](#)”.

- 2 Open your Internet browser and enter the URL for iManager.

The URL is `http://server_ip_address/nps/imanager.html`. Replace *server_ip_address* with the IP address or DNS name of a server in the cluster or with the IP address for Apache-based services.

- 3 Enter your username and password.
- 4 In the left column, locate and click *Clusters > Cluster Options*. You can configure and manage cluster using this section.
- 5 Specify a cluster name or browse and select one, the page displays cluster resources for the selected cluster.
- 6 Select and click the cluster object, that was created in [Section 10.3.1, “Creating NSS Shared Disk Partition,” on page 88](#) on the Archive volume.
- 7 On the Cluster Resources Properties page, select the *Scripts* tab.

7a Click *Load Script*.

The *Load Script* displays the script generated during creation of the shared NSS pool.

- 7b** You can either copy the contents of the Archive load script from “[Archive Load Script Configuration](#)” on page 92 or from the load script of the *AV_Template* and paste it in the *Load Script* of the cluster object just before the exit 0 statement.

7c Click *Unload Script*.

The *Unload Script* displays the script generated during creation of the shared NSS pool.

- 7d** You can either copy the contents of the Archive unload script from [“Unload Script Configuration” on page 93](#) or from the unload script of the [AV_Template](#) and paste it in the Unload Script just before the exit 0 statement.
- 7e** Click *Monitor Script*.
The monitor script displays the script to monitor the Archive service.
- 7f** You can copy the contents of the Archive monitor script from the [“Monitor Script Configuration” on page 95](#) and paste it in the *Monitor Script*.
- 7g** Click OK.
- 8** Using iManager, configure the Archive Server. For more information, see [“Configuring Archive Volume” on page 54](#).
 - 8a** Click *Archive Versioning > Archive Server Properties* to open the Archive Server Properties window in the main panel.
 - 8b** Specify or browse the name of the eDirectory™ context for the cluster enabled archive server you want to manage.
 - 8c** In the Volume Path field, specify the shared NSS volume and path on the archive server to store the versions of files.
 - 8d** Click OK.
- 9** Continue with [Section 10.5, “Configuring Jobs using iManager,” on page 95](#)

10.4.3 Archive Load, Unload and Monitor Scripts

Archive Load Script Configuration

The Archive load script page should already be displayed. The load script contains commands to start the Archive service. You must customize some commands for your specific Archive configuration.

The load script appears similar to the following example:

```
#!/bin/bash
./opt/novell/ncs/lib/ncsfuns

# define the IP address
RESOURCE_IP=a.b.c.d

# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=name
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/ark
# define the mount point
MOUNT_POINT=/mnt/ark

# activate the volume group
exit_on_error vgchange -a ey $VOLGROUP_NAME

# mount the file system
exit_on_error mount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS
```

```

# add the IP address
# exit_on_error add_secondary_ipaddress $RESOURCE_IP

# delete existing PostgreSQL pid file
# <data_dir> is the value of DATA_PATH variable from
# /opt/novell/arkmanager/bin/pg_init.sh.
# DATA_PATH must point to a location on shared disk
# define DATA_DIR
DATA_DIR=path_to_data_dir

ignore_error rm /<$DATA_DIR>/postmaster.pid

# copy state file
ignore_error mv /opt/novell/arkmanager/arkState.xml /opt/novell/arkmanager/
arkState.xml.bak
ignore_error mv /<$DATA_DIR>/arkState.xml /opt/novell/arkmanager

# start Archive Version Services
exit_on_error /etc/init.d/novell-ark start

# return status
exit 0

```

Configuring Archive Load script

To customize the Archive and Versioning load script for your specific configuration:

- 1 Edit the following line to assign a unique IP address to Archive and Versioning cluster resource:

```
RESOURCE_IP=a.b.c.d
```

Replace *a.b.c.d* with the IP address you want to assign to the AV cluster resource.

The IP address for the AV cluster resource allows clients to reconnect to that address regardless of which server is hosting it.

- 2 Edit the value of MOUNT_DEV to the location of the device created on configuring the shared Linux POSIX Volume.
- 3 Edit the value of MOUNT_POINT to the location you want to mount the device specified by [MOUNT_DEV](#).
- 4 Edit the value of the DATA_PATH variable to assign your specific data directory.

```
ignore_error rm /<data_dir>/postmaster.pid
```

- 5 Click Next and continue with the unload script configuration.

Unload Script Configuration

The Archive unload script page should already be displayed. The unload script contains commands to stop the Archive service. You must customize some commands for your specific Archive configuration.

The unload script appears similar to the following example:

```
#!/bin/bash
./opt/novell/ncs/lib/ncsfuncs

# define the IP address
RESOURCE_IP=a.b.c.d

# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=name
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/ark
# define the mount point
MOUNT_POINT=/mnt/ark

# define the DATA_DIR
# <data_dir> is the value of DATA_PATH variable from
# /opt/novell/arkmanager/bin/pg_init.sh.
# DATA_PATH must point to a location on shared disk
DATA_DIR=/mnt/a/b/c/d

# stop Archive Version Services
ignore_error /etc/init.d/novell-ark stop

# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# Stop the postgres sql database
# <arkuser> is PostgreSQL user name
ignore_error sudo -u <arkuser> pg_ctl stop -D $DATA_DIR -m fast

# delete existing PostgreSQL pid file
ignore_error rm $DATA_DIR/postmaster.pid

# copy state file
ignore_error cp /opt/novell/arkmanager/arkState.xml $DATA_DIR/arkState.xml

# umount the file system
sleep 10 # if not using SMS for backup, please comment out this line
exit_on_error umount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

#deactivate the volume group
exit_on_error vgchange -a n $VOLGROUP_NAME

# return status
exit 0
```

Configuring Archive Unload script

To customize the Archive and Versioning unload script for your specific configuration:

- 1 Edit the following line to assign a unique IP address to Archive and Versioning cluster resource:

```
RESOURCE_IP=a.b.c.d
```

Replace *a.b.c.d* with the same IP address you specified in the load script.

- 2 Edit the value of MOUNT_DEV to the location of the device created on configuring the shared Linux POSIX Volume.
- 3 Edit the value of MOUNT_POINT to the location of the mounted path of the Linux POSIX Volumes on a shared disk.
- 4 Edit the value of the DATA_PATH variable to assign your specific data directory, then click next.

```
DATA_DIR=/mnt/a/b/c/d
```

Monitor Script Configuration

The monitor script contains commands to monitor the Archive service. You must customize some commands for your specific Archive configuration.

The unload script appears similar to the following example:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuns

# define the IP address
RESOURCE_IP=a.b.c.d

# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=name
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/ark
# define the mount point
MOUNT_POINT=/mnt/ark

# check the logical volume
exit_on_error status_lv $MOUNT_DEV

# check the file system
exit_on_error status_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# check the IP address
exit_on_error status_secondary_ipaddress $RESOURCE_IP

# check AV
exit_on_error /etc/init.d/novell-ark status

# return status
exit 0
```

10.5 Configuring Jobs using iManager

To set up jobs for Archive and Version Services on each node, specify the job properties like job name, source server configuration properties, delete policy and run schedule. You need to start configuring with the master node and then go in rotation making other nodes master, one by one.

For each server in the cluster, you can configure jobs using iManager plug-in. For more information, see [Chapter 7, “Configuring Jobs in iManager,” on page 53](#)

NOTE: If you modify or create the versioning job, you must copy the `/etc/opt/novell/arkmanager/conf/arkConfig.xml` file to all of the archive servers in the cluster.

For information on editing `arkConfig.xml`, see [“Sample Configuration for arkConfig.xml” on page 117](#)

10.6 Managing Archive Resource

For information on managing Cluster Resource for Archive and Version Services 2.1, see [OES 11: Novell Cluster Services 2.0 for Linux Administration Guide](#).

11 Migrating Novell Archive and Version Services to OES 11

The OES 11 Migration Tool has a plug-in architecture and is made up of Linux command line utilities with a GUI wrapper. You can migrate Archive and Version Services to an OES 11 server through either the GUI Migration Tool or through the command line utilities.

To get started with migration, see *Using the Migration Tool GUI* in the [OES 11: Migration Tool Administration Guide](#).

For more information on migrating Novell Archive and Version Services, see “[Migrating Novell Archive and Version Services to OES 11](#)” in the [OES 11: Migration Tool Administration Guide](#).

12 Security Considerations for Archive and Version Services

This section discusses the following security issues for Novell Archive and Version Services:

- ♦ [Section 12.1, “NSS Encrypted Volumes,” on page 99](#)
- ♦ [Section 12.2, “Secure Data Transfer,” on page 99](#)
- ♦ [Section 12.3, “eDirectory and PostgreSQL Administrator Passwords,” on page 99](#)
- ♦ [Section 12.4, “User Authentication,” on page 100](#)

12.1 NSS Encrypted Volumes

If you plan to version files from one or more encrypted source volumes, the archive volume must also be an encrypted volume. Otherwise, the data is stored nonencrypted on the archive volume.

For information about NSS Encrypted Volume Support, see Managing Encrypted NSS Volumes in the *OES 11: NSS File System Administration Guide for Linux*.

The data is not secure when transported between the source volume and the archive volume. For information, see [Section 12.2, “Secure Data Transfer,” on page 99](#).

12.2 Secure Data Transfer

Because the archive server moves data in nonencrypted format, the archive server should be located behind the corporate firewall. If you are versioning data on remote servers, use a virtual private network (VPN) connection between the two.

12.3 eDirectory and PostgreSQL Administrator Passwords

The passwords for ArkManager administrator user in Novell eDirectory and PostgreSQL are passed using YaST and stored in CASA.

To modify the password for the ArkManager administrator user in eDirectory or the password for the database administrator user in PostgreSQL, see [Section 6.3, “Updating the Archive Server Configuration,” on page 50](#).

12.4 User Authentication

Archive and Version Services uses eDirectory to authenticate users who access the archive database. For more information, see [Section 4.1, “Network Architecture Prerequisites and Guidelines,” on page 33](#).

In the following scenarios the user can restore or view all the previous versions of the files or directories:

- ♦ A user assigned with write permissions can view all the previous versions of the files or directories, even if he or she was previously not allowed to access the older versions of that file.
- ♦ A user assigned with read permissions can view all the previous versions of the file even if he has no write permission for the current location. The user can restore the files or directories to a different location where he or she has write permissions.

A XML Elements and Attributes for ArkConfig

This section defines the XML elements and attributes that you use to configure versioning jobs for your Novell Archive and Version Services 2.1 for OES 11 server. The XML elements have a particular hierarchy that you must keep in mind as you define the Basic element, Defaults element, and multiple Job elements. See [Figure A-1 on page 102](#) to understand the parent-child relationships between XML elements defined for the `etc/opt/novell/arkManager/conf/arkConfig.xml` file.

The following table defines the XML elements and attributes that you use in the `/var/opt/novell/arkmanager/arkConfig.xml` file. Some elements appear first as children of elements in the next-higher level of the hierarchy and as parents to their own set of child elements in the next-lower level of the hierarchy.

The examples provided for the elements are sample code to illustrate how the configured properties appear in the XML file. You must modify the properties to meet your specific needs as you create your `etc/opt/novell/arkManager/conf/arkConfig.xml` file.

For a sample of the basic configuration, see [“Sample Configuration for arkConfig.xml” on page 117](#).

NOTE: The user credentials for edirectory and PostgreSQL database are not authenticated using `arkConfig.xml`. For security reasons, the credentials are passed using YaST and stored in CASA store.

Figure A-1 Hierarchical Parent-Child Relationships between XML Elements Used in the arkConfig.xml File

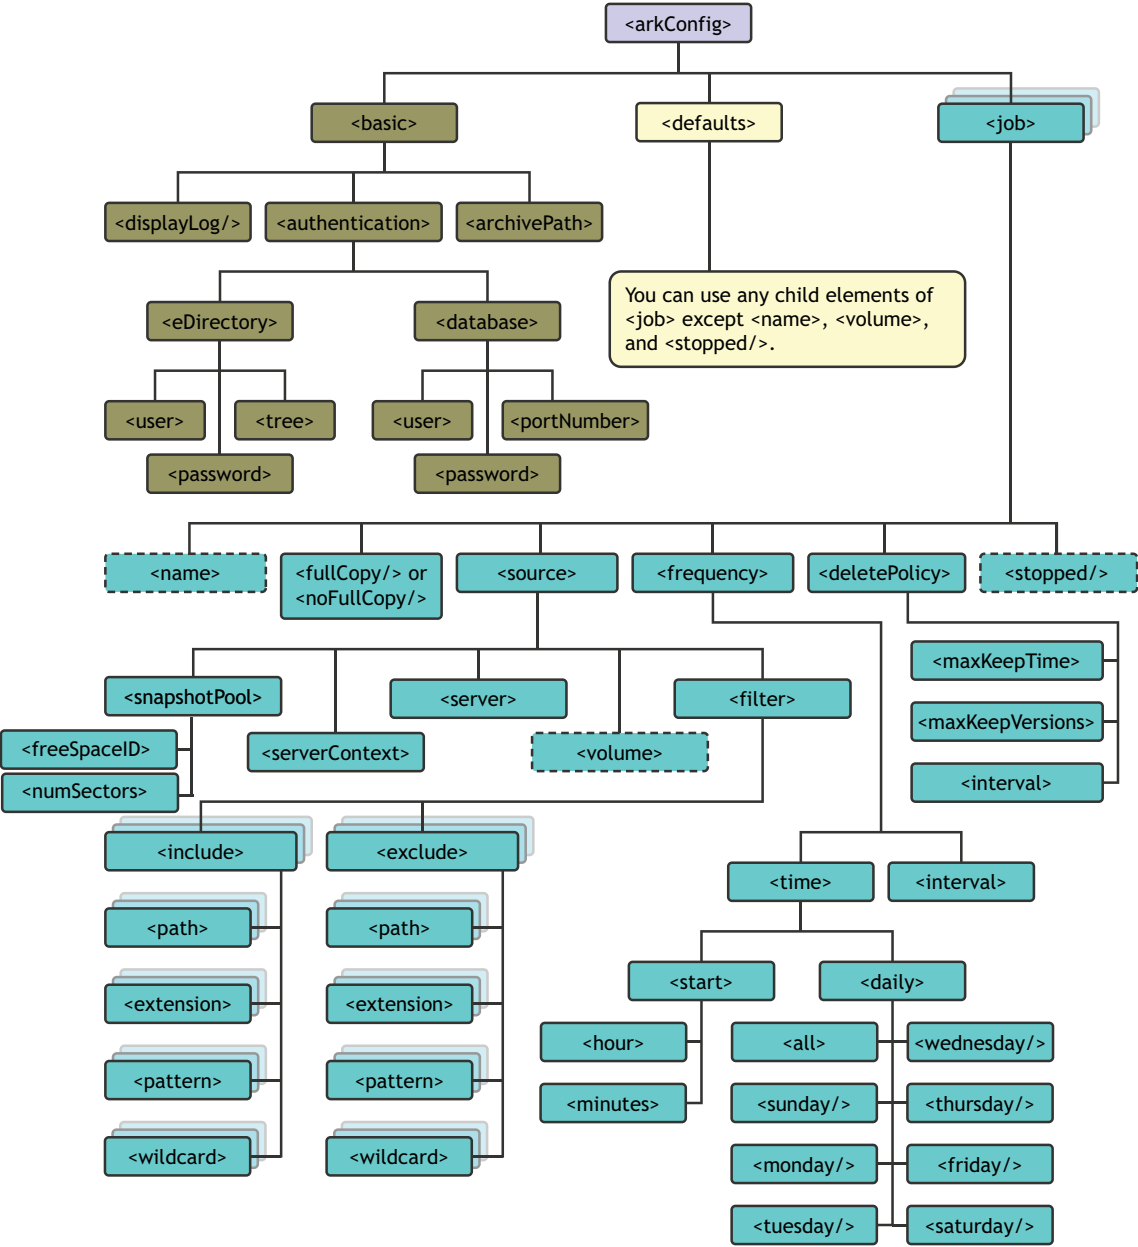


Table A-1 Elements of the *arkConfig.xml* File

Parent Element	Child Elements and Attributes	Description
<arkConfig> (Root element)	<basic>	<p>A single Basic tag set surrounds the management information that applies to all jobs. Its child elements specify where to display the log, the authentication information, and the storage location where the archive data (versions) reside.</p> <p>At a minimum, configure the <authentication> information for ArkManager 2.0 to run.</p> <p>For information about its child elements, see <basic> in the Parent Element column of this table.</p>
	<defaults>	<p>A single Defaults tag set surrounds the child elements that define the default settings to use if the elements are not set in individual jobs.</p> <p>Each job uses the values configured in the Defaults element only for those values that are not otherwise set for that job.</p> <p>For information about its child elements, see <defaults> in the Parent Element column of this table.</p>
	<job>	<p>A Job tag set surrounds the child elements that define the settings for a given job. Create one job for each unique information set that you want to version. An archive server can version a given source volume with only one job.</p> <p>The values configured in a given job override the default settings only for that job. If a value is not set in the Job element, the default settings apply for only that element.</p> <p>For information about its child elements, see <job> in the Parent Element column of this table.</p>

Parent Element	Child Elements and Attributes	Description
<code><defaults></code> (Child to <code><arkConfig></code>)	<code><fullCopy/></code> <code><noFullCopy/></code> <code><source></code> <code><frequency></code> <code><deletePolicy></code>	<p>For information about these child elements, see the Job element (<code><job></code>) in this table.</p> <p>The Defaults element can contain any tags that you use in a Job element, except for <code><name></code>, <code><volume></code>, and <code><stopped/></code>.</p> <p>Although you can have multiple Job elements, only a single Defaults element appears in the <code>arkConfig.xml</code> file.</p>

Parent Element	Child Elements and Attributes	Description
<job> (Child to <arkConfig>)	<name>	<p>Specifies the unique name of this versioning job. This name can contain spaces.</p> <p>You might want the job name to reflect information about the data that is being versioned. For example, if you have a job that versions data for a <code>user1</code> volume on the <code>svr1</code> server, you might name the job:</p> <pre><name>user1[svr1]</name></pre> <p>Use the <code><name></code> tag as a child only of the Job element; it is invalid in the Defaults element.</p>
	<fullCopy/>	<p>(Optional) If used, sets a flag so that the job copies all files that pass the filter from the source volume to the archive database the first time the job runs or if a full copy runs but did not finish previously.</p> <p>In any given Job element or Defaults element, use <code><fullCopy/></code> or use <code><noFullCopy/></code>, not both. If neither one is specified, <code><fullCopy/></code> is the default software setting.</p> <p>This is an empty element; it has no child elements and no data values.</p>
	<noFullCopy/>	<p>(Optional) If used, sets a flag so that the job does not copy all files that pass the filter from the source volume to the archive database the first time the job runs.</p> <p>This is an empty element; it has no child elements and no data values.</p>
	<source>	<p>Contains child elements that define the specific server context, server, paths, directories, or file types to be archived for this job.</p> <p>For information about its child elements, see <source> in the Parent Element column of this table.</p> <p>NOTE: If a job is configured using IP address or hostname in the <code><source></code> element, ensure that the IP address or hostname is same as the IP address or hostname of the server where the data to be versioned is located.</p>
	<frequency>	<p>Contains child elements that define the start time and intervals for versioning in this job.</p> <p>For information about its child elements, see <frequency> in the Parent Element column of this table.</p>

Parent Element	Child Elements and Attributes	Description
<job> (cont'd)	<deletePolicy>	Contains child elements that specify when to delete a version, such as by age of the version (elapsed time of creation), by the number of versions that exist (oldest deleted first), or by the interval of time before the next versioning process. For information about its child elements, see <deletePolicy> in the Parent Element column of this table.
	<stopped/>	Specify this property to define the job but leave it in a Stopped state until you manually activate the job. If the Stopped parameter is not used, the job starts, according to the Run Schedule settings, the next time ArkManager runs. Use the <stopped/> tag as a child only of the Job element; it is invalid in the Defaults element.
<authentication> (Child to <basic>)	<eDirectory>	Contains child elements that specify the Novell eDirectory authentication information, including the administrator user, the password, and the eDirectory tree where the administrator user, archive server, and source servers reside. For information about its child elements, see <eDirectory> in the Parent Element column of this table.
	<database>	Contains child elements that contain information about the archive server's ArkSQL database. This user must have server-specific rights to the ArkSQL server. For information about its child elements, see <database> in the Parent Element column of this table.
<eDirectory> (Child to <authentication>)	<user>	This attribute is not used for user authentication.
	<password>	This attribute is not used for user authentication.
	<tree>	This attribute is not used for user authentication. NOTE: The eDirectory credentials are passed using YaST and are stored in the CASA store.
<database> (Child of <authentication>)	<user>	This attribute is not used for user authentication.
	<password>	This attribute is not used for user authentication.
	<portNumber>	This attribute is not used for user authentication. NOTE: The PostgreSQL credentials are passed using YaST and are stored in the CASA store.

Parent Element	Child Elements and Attributes	Description
<source> (Child to <defaults> and to <job>)	<snapshotPool>	<p>Specifies the name of the destination pool where a snapshot of the source volume is created temporarily at the beginning of each versioning. For example:</p> <pre><snapshotPool>pool_name </snapshotPool></pre> <p>Snapshots make it possible to save point-in-time versions of all eligible source files, even if a file is open at the time. Eligible files are those files that exist at the end of the epoch, changed during the epoch, and met the filter requirements.</p> <p>By default, if no snapshot pool is specified, ArkManager copies the eligible source files directly from the source volume. Exclusively open files cannot be versioned and data might be inconsistent.</p> <p>If the specified snapshot pool is inactive or otherwise not available when a versioning job begins, the job copies the files directly from the source volume.</p>
	<freeSpaceID>	ID of the free space object to be used for storing snapshot data. For example, hda_freespace1 or hda.
	<numSectors>	Specifies the number of sectors on free space to be used for storing snapshot data. The value needs to be an integer, for example, 419430400.
	<serverContext>	<p>Specifies the unique name of the context of the server that hosts the source volume, which contains the data that the job versions and saves to the archive database.</p> <p>Type the server context in the Novell common dot format, from lowest to highest level. It does not include the tree. For example:</p> <pre><serverContext> grp1.dept1.site1.examplecompany </serverContext></pre> <pre><serverContext> sales.mktg.eur.acme </serverContext></pre>
	<server>	<p>Specifies the unique name of the server in the specified context where the original data is stored. For example, to specify a server named <code>srv1</code>:</p> <pre><server>srv1</server></pre> <p>Use the <server> tag as a child only of the Job element; it is invalid in the Defaults element.</p>

Parent Element	Child Elements and Attributes	Description
<source> (cont'd)	<volume>	<p>Specifies the unique name of the OES 11 NSS volume in the specified source server where the original data is stored.</p> <p>For example, to specify <code>user1</code> as the source volume:</p> <pre><volume>user1</volume></pre> <p>Do not place a colon after the name of the source volume.</p> <p>Use the <volume> tag as a child only of the Job element; it is invalid in the Defaults element.</p>
	<filter> <pre><filter overrideDefaults="false"></pre> <pre><filter overrideDefaults="true"></pre>	<p>(Optional) Contains child elements and attributes that specify the criteria for filtering data so that only selected data is versioned.</p> <p>Use the attributes optionally to indicate whether the filter is to be used in combination with filters set in the Defaults element (false) or if the filter is to be used in place of the filters set in the Defaults element (true).</p> <p>For information about its child elements, see <filter> in the Parent Element column of this table.</p>
<filter> (Child to <source>)	<include>	<p>(Optional) Contains child elements that specify traits of data you want to include in the job.</p> <p>For information about its child elements, see <include> in the Parent Element column of this table.</p>
	<exclude>	<p>(Optional) Contains child elements that specify traits of data you want to exclude in the job.</p> <p>For information about its child elements, see <exclude> in the Parent Element column of this table.</p>

Parent Element	Child Elements and Attributes	Description
<code><include></code> (Child to <code><filter></code>)	<code><path></code>	<p>(Optional) Specifies the relative path of directories in the specified volume that you want to include in the versioning process. If you want to include eligible files only in the specified paths, make sure to exclude the root path in an exclude path statement.</p> <p>For example, to specify a relative path of the <code>volume_name:\dept1\data</code> directory that contains data you want to include:</p> <pre><include> <path>\dept1\data\</path> </include></pre> <p>Repeat this element for each directory path in the specified volume that contains data that you want to version.</p>
	<code><extension></code>	<p>(Optional) Specifies the extension of files in the specified volume that you want to include in the versioning process. Use the preceding dot followed by the characters of the file extension.</p> <p>Repeat this element for each file extension type that you want to version.</p> <p>For example, to specify that you want to version only files with no extension (.) and files with .xxx and .yyy extensions:</p> <pre><include> <extension>.</extension> <!-- No extension --> <extension>.xxx</extension> <extension>.yyy</extension> </include></pre>
	<code><pattern></code>	<p>(Optional) Specifies the regular expression pattern to match for files that you want to include in the versioning process.</p> <p>For example, to set criteria to include only files with names that start with the letter "a".</p> <pre><include> <pattern>.*\a.*</pattern> </include></pre> <p>Repeat this element for each pattern that you want to match to identify files for versioning.</p> <p>This element does not support the full set of PERL regular expressions. For more information, see "Pattern Elements" on page 30.</p>

Parent Element	Child Elements and Attributes	Description
<code><include></code> (cont'd)	<code><wildcard></code>	<p>(Optional) Functions like a wildcard in directory searches. Replace characters with an asterisk (*) to search for files that match. For example, to include files that start with d of type .sxi:</p> <pre> <include> <wildcard>d*sxi</wildcard> </include> </pre>
<code><exclude></code> (Child to <code><filter></code>)	<code><path></code>	<p>(Optional) Specifies the relative path of directories in the specified volume that you want to exclude from the versioning process.</p> <p>For example, to specify a relative path of the volume_name:\dept1\apps directory that contains application data you want to exclude from versioning:</p> <pre> <exclude> <path>\dept1\apps</path> </exclude> </pre> <p>Repeat this element for each directory path in the specified volume that contains data that you do not want to version.</p>
	<code><extension></code>	<p>(Optional) Specifies the extension of files in the specified volume that you want to exclude from the versioning process. Use the preceding dot followed by the characters of the file extension.</p> <p>Repeat this element for each file extension type that you do not want to version.</p> <p>For example, to specify that you want to exclude .zzz files from versioning:</p> <pre> <exclude> <extension>.zzz</extension> </exclude> </pre>

Parent Element	Child Elements and Attributes	Description
<code><exclude></code> (cont'd)	<code><pattern></code>	<p>(Optional) Specifies the regular expression pattern to match for files that you want to exclude from the versioning process.</p> <p>For example, to set criteria to exclude files with names that start with the letter “a”:</p> <pre><exclude> <pattern>.*\a.*</pattern> </exclude></pre> <p>Repeat this element for each pattern that you want to match to identify files to be excluded from versioning.</p> <p>This element does not support the full set of PERL regular expressions. For more information, see “Pattern Elements” on page 30.</p>
	<code><wildcard></code>	<p>(Optional) Functions like a wildcard in directory searches. Replace characters with an asterisk (*) to search for files that match. For example, to exclude files that start with d of type .tmp:</p> <pre><exclude> <wildcard>d*tmp</wildcard> </exclude></pre>

Parent Element	Child Elements and Attributes	Description
<frequency> (Child to <defaults> and to <job>)	<time>	<p>(Conditional) Contains child elements that specify the start time, based on a 24-hour local time scale, of jobs that are scheduled to occur on specified days of the week. A valid hour entry is an integer value ranging from 0 to 23.</p> <p>Use either the Time element or Interval element, but not both in the same Job element.</p> <p>For example, to start the versioning at 11:15 p.m. (23:45) every day of the week:</p> <pre> <frequency> <time> <start> <hour>23</hour> <minute>15</minute> </start> <daily> <all/> </daily> </time> </frequency> </pre> <p>For information about its child elements, see <time> in the Parent Element column of this table.</p>
	<interval> <interval unit="seconds"> <interval unit="minutes"> <interval unit="hours"> <interval unit="days">	<p>(Conditional) Specifies the elapsed time between scheduled versions in seconds, minutes, hours, or days.</p> <p>You must use one of the attributes in the Interval tag to specify the units of the integer value you specify as the interval of time to wait between when the versioning begins. For example, to specify an interval of 30 minutes:</p> <pre> <interval unit="minutes">30 </interval> </pre> <p>You must specify a non-zero value for the interval or the job cannot start. If no interval value is set, 0 (zero) is the default setting.</p> <p>If no interval unit is set, seconds are the default unit setting.</p> <p>Use either the Time element or Interval element, but not both in the same Job element.</p>

Parent Element	Child Elements and Attributes	Description
<time> (Child to <frequency>)	<start>	<p>Contains child elements that specify the hour and minutes that the job is scheduled to run on the specified days of the week.</p> <p>For information about its child elements, see <start> in the Parent Element column of this table.</p>
	<daily>	<p>Contains child elements that specify the days of the week that the job is scheduled to run. You must specify at least one day or the job cannot start.</p> <p>For information about its child elements, see <daily> in the Parent Element column of this table.</p>
<start> (Child to <time>)	<hour>	<p>Specifies the hour of the day that the job begins, based on a 24-hour clock. A valid hour entry is an integer value ranging from 0 to 23. For example, to set the hour to 10 p.m.:</p> <pre><hour>22</hour></pre> <p>If no hour is set, midnight (00) is the default hour setting.</p>
	<minute>	<p>Specifies the minutes of the hour in the day that the job is scheduled to begin. A valid minutes entry is an integer value ranging from 0 to 59. For example, to set the minutes after the hour to 45 minutes:</p> <pre><minute>45</minute></pre> <p>If no minute value is set, zero (00) is the default minute setting.</p>
<daily> (Child to <time>)	<monday/> <tuesday/> <wednesday/> <thursday/> <friday/> <saturday/> <sunday/> <all/>	<p>Specify one or more of the child elements as the days of the week you want to schedule the versioning to occur. If no <daily> value is set, the job cannot start.</p> <p>For example, to specify versioning to occur on Tuesday, Friday, and Sunday of each week:</p> <pre><daily> <tuesday/> <friday/> <sunday/> </daily></pre> <p>Use the <all/> tag to specify all seven days of the week. For example:</p> <pre><daily> <all/> </daily></pre> <p>Each of the <daily> child elements is an empty element; it has no child elements and no data values.</p>

Parent Element	Child Elements and Attributes	Description
<code><deletePolicy></code> (Child to <code><defaults></code> and to <code><job></code>)	<code><maxKeepTime></code> <code><maxKeepTime unit="seconds"></code> <code><maxKeepTime unit="minutes"></code> <code><maxKeepTime unit="hours"></code> <code><maxKeepTime unit="days"></code> <code><maxKeepTime keepCurrentCopy="true"></code> <code><maxKeepTime keepCurrentCopy="false"></code>	<p>Use one of the unit attribute values in the <code><maxKeepTime></code> tag to specify the units of the integer value you specify as maximum lifetime of a version.</p> <p>If no <code>maxKeepTime</code> value is set, file versions are retained indefinitely.</p> <p>If a value is specified without a unit attribute, "seconds" is the default unit of the value specified.</p> <p>Use the <code>keepCurrentCopy</code> attribute to designate that at least one file version of current files remains in the database, even if the <code>maxKeepTime</code> elapses.</p> <p>If <code>keepCurrentCopy</code> is set to True (default), the archive keeps an existing file version as long as its source file is current on the source volume, beyond the <code>maxKeepTime</code>. After the user deletes the current source file, the deletion is noted at the next scheduled epoch. If the file version's age is within the <code>maxKeepTime</code>, the archive database retains a copy of the file version until its <code>maxKeepTime</code> elapses. When the file version's age exceeds the <code>maxKeepTime</code>, the archive deletes the file version at the next scheduled delete time.</p> <p>If <code>keepCurrentCopy</code> is set to False, the archive deletes a file version when it exceeds the <code>maxKeepTime</code>, even if the file version is the only one in the archive database for a given source file, and whether the source file is current or deleted.</p> <p>The <code>keepCurrentCopy</code> attribute is optional. If the <code>keepCurrentCopy</code> attribute is not otherwise specified, the default value is True.</p> <p>For example, to keep versions for 90 days before purging, to keep at least the most current version, and to schedule the purging of versions with 1-hour intervals:</p> <pre> <deletePolicy> <maxKeepTime unit="days" keepCurrentCopy="true">90 </maxKeepTime> <interval unit="hours">1 </interval> </deletePolicy> </pre>

B Sample Configuration Files

This section illustrates the use of XML file used to configure Novell Archive and Version Services 2.1 for OES 2. For your convenience, all sample files can be found in the `etc/opt/novell/arkManager/conf` directory of your archive server. Use the following table to determine which samples to use.

Table B-1 Possible Tasks

If You Need To	View This Sample File	Sample File
The sample files for Novell Archive and Version Service 2.1 for OES 11 is located in the <code>etc/opt/novell/arkmanager/conf</code> directory		
Configure all properties for the archive server, individual jobs, and default job settings in the <code>arkConfig.xml</code> file.	“arkConfig.xml - Sample Configuration file” on page 117	<code>arkconfig.xml</code>

B.1 arkConfig.xml - Sample Configuration file

To configure versioning jobs, modify the `arkConfig.xml` located at `/etc/opt/novell/arkmanager/conf/` as per the sample configuration file.

Sample Configuration for arkConfig.xml

```
<arkConfig>
  <basic>
    <displayLog/>
    <authentication>
      <eDirectory>
        <user>edirectory admin user(eg: cn=admin.o=abc)</user>
        <password>edirectory admin user's password</password>
        <tree>ABC-TREE</tree>
      </eDirectory>
      <database>
        <user>PostGRESQL db user for ark server</user>
        <password>PostGRESQL db user</password>
        <portNumber>portnumber on which PgSQL is running</portNumber>
      </database>
    </authentication>
    <archivePath>NSSARKVOL:/archive_directory</archivePath>
  </basic>
  <job>
    <name><![CDATA[ job1 ]]></name>
    <noFullCopy/>
    <frequency>
      <interval unit="minutes">15</interval>
    </frequency>
    <source>
      <serverContext><![CDATA[novell]]></serverContext>
```

```

        <server><![CDATA[Archive-server.novell]]></server>
        <volume><![CDATA[primevol2]]></volume>
        <snapshotPool><![CDATA[primepool2]]></snapshotPool>
        <freeSpaceID><![CDATA[hda_freespace1]]></freeSpaceID>
        <numSectors><![CDATA[419430400]]></numSectors>
        <filter overrideDefault="false">
            <include><extension><![CDATA[.txt]]></extension></include>
        </filter>
    </source>
    <deletePolicy/>
</job>
<job>
    <name><![CDATA[job2]]></name>
    <frequency>
        <interval unit="minutes">20</interval>
    </frequency>
    <source>
        <serverContext><![CDATA[abc]]></serverContext>
        <server><![CDATA[SERVER_NAME]]></server>
        <volume><![CDATA[NSSVOL2_NAME]]></volume>
        <snapshotPool/>
        <filter overrideDefault="false">
            <include><extension><![CDATA[.tmp]]></extension></include>
            <exclude><path><![CDATA[\\test1\\]]></path></exclude>
            <exclude><extension><![CDATA[.doc]]></extension></exclude>
            <exclude><wildcard><![CDATA[imp*.exe]]></wildcard></exclude>
        </filter>
    </source>
    <deletePolicy/>
</job>
</arkConfig>

```

C Troubleshooting

This section discusses potential issues and workarounds for Novell Archive and Version Services.

- ♦ [Section C.1, “Archive Server,” on page 119](#)
- ♦ [Section C.2, “Migration Issues,” on page 123](#)
- ♦ [Section C.3, “NSS File Version Utility,” on page 123](#)

C.1 Archive Server

- ♦ [“Job fails to start” on page 119](#)
- ♦ [“Undefined Job Name” on page 120](#)
- ♦ [“Versions are not available” on page 120](#)
- ♦ [“Archive server on OES 2 SP1 NetWare server fails to version files stored on Netware NSS volume” on page 120](#)
- ♦ [“Archive Server fails to start” on page 121](#)
- ♦ [“Archive Server fails to start. Ensure that root is part of the novlxtier group” on page 121](#)
- ♦ [“I am unable to view versions using vi,vim,gedit,kwrite?” on page 121](#)
- ♦ [“Unable to get versions of the symbolic file” on page 122](#)
- ♦ [“On the stopping the PostgreSQL instance of Archive service, the database status still displays running.” on page 122](#)
- ♦ [“The database port change is not reflected on modifying arkConfig.xml file.” on page 122](#)
- ♦ [“Browse icon not functional for restoring files with NetStorage in a Firefox browser” on page 122](#)

Job fails to start

Using iManager, job fails to run with an error, “Error authenticating connection, check if eDirectory user name and password are correct.”

Possible Cause: The eDirectory user is not LUM enabled and part of novlxtier LUM group on the OES 2 server running Archive and Version Services.

Action: Ensure that the eDirectory user is LUM enabled on the server running Archive and Version Services.

- 1 For more information on LUM enabling the eDirectory user, see the [OES 11: Novell Linux User Management Administration Guide](#).
- 2 Add the eDirectory user to the “novelxtier” LUM group. The group is created on installing the OES 2 server.

3 To refresh the LUM cache, run `namconfig cache_refresh`.

4 Start the Archive and Version Service.

```
rcnovell-ark start
```

On restarting the service job runs as per the schedule.

Possible Cause: If a job is manually configured in the `arkConfig.xml` file using IP address or hostname of the source server.

Action: Ensure that the IP address or hostname mentioned in the `source` element of the `arkConfig.xml` is same as the IP address or hostname of the server, where the data to be versioned is located. On updating the IP address or hostname in the `arkConfig.xml` file, restart the Archive server.

Possible Cause: Casa credentials are not set correct

Action: To reset the password in the CASA store run the following command:

```
/opt/novell/arkmanager/bin/arkReconfig.sh --edir-admin=<adminusername>
```

You will be prompted for the password.

Restart the arkserver using `rcnovell-ark restart`

Undefined Job Name

In iManager, *Archive Versioning>Archive Jobs*, name of the job is listed as *Undefined Job Name* under the *Jobs* section.

Possible Cause: In the Job settings page, source server name is without context.

Action: Delete the job named *Undefined Job Name*. Create a new job and ensure that the source server name is with context. For example, *Archive-server.org*

Versions are not available

Possible cause: The user does not have access to Unix workstation or is not LUM-enabled.

Action: The eDirectory user needs to be LUM-enabled and given access to the Unix workstation by the administrator. For more information, see the [Steps to LUM-enable the user](#).

Possible cause: The Archive server and Client desktop workstation system clocks are not synchronized.

Action: Synchronize the Archive and Client desktop system time.

Archive server on OES 2 SP1 NetWare server fails to version files stored on Netware NSS volume

Possible cause: The `archive.cfg` file is present at the root level on the NSS Volume of OES 2 SP1 NetWare server.

Action: To get versions, delete the `archive.cfg` file.

Archive Server fails to start

View the server log files available at `/var/opt/novell/arkmanager/logs`.

Possible cause: NSS is not running

Action: Use `rcnovell-nss status` to check if NSS is running. If it is not running, start NSS with the `rcnovell-nss start` command.

Possible cause: Adminfs is not running

Action: Check if adminfs is running with `rcadminfs status`. If it is not running, start adminfs with the `rcadminfs start`

Possible cause: Postgresql driver not found.

Action: Check if the `postgresql-8.1-404.jdbc3.jar` file is available in `/usr/share/pgsql` directory. If the file does not exist, you need to install the `postgresql-jdbc.8.1.rpm` file.

Possible cause: One of the cause could be the version of the JDBC driver. By default, the archive server functions with `postgresql-8.1-404.jdbc2ee.jar` file installed with Novell Open Enterprise Server x.x Product CD.

Action: By default, the archive server functions with `postgresql-8.1-404.jdbc2ee.jar` file installed with Novell Open Enterprise Server 2 Product CD. If you have installed any other version of `.jar` file on your system, you need to edit the classpath available in the `novell-ark` file.

For example, if you have installed the JDBC 3 driver and version of respective jar file is `postgresql-8.2-504.jdbc3.jar`. Edit the classpath in the `novell-ark` file as follows:

```
CP_PATH5: /usr/share/pgsql/postgresql-8.2-504.jdbc3.jar
```

Possible cause: PostgreSQL database is not running.

Action: Delete the `postmaster.pid` file from the PostgreSQL database data path, then restart Archive server.

Archive Server fails to start. Ensure that root is part of the novlxtier group

Cause: You are running the Archive server as a `root` user.

Action: Reconfigure Archive and Version Services.

- 1 Reconfigure Archive and Version Service to run as a non-root user.
- 2 Add the non-root user to the "novlxtier" LUM group. The group is created on installing the OES 2 SP1 or later server.
- 3 Start Archive and Version Service
`rcnovell-ark start`

I am unable to view versions using vi,vim,gedit,kwrite?

Possible cause: The text editor for example, vi, creates a swap file with a `.swp` extension on editing a file. The creation of the swap files avoids displaying versions of original files. As per the functionality of text editor, it creates temporary files during the process of modification and saves it with original name after successful modification. The Archive server is designed to capture changes to the file and create versions. So all the changes are tracked and versions are stored even for the temporary files.

Action: Create an *Include* or *Exclude* filter to avoid versioning of temporary files like .swp or .temp. You can set filter using iManager, *Archive Versioning* > *Archive Server Properties* > *Source Server information* > *Filter*. Set a filter and specify the extension of files that you want to exclude in the *Extension* field. For example, .tmp

Unable to get versions of the symbolic file

Possible cause: The original file is not stored on NSS file system configured for Archive and Version Services.

Explanation: For example, if original file is stored on ext3, reiserfs or any other file system and symbolic link is created on NSS Primary volume. Any modification to the original file will not create versions on the Primary volume. Versions are available of files stored on NSS file system which is configured for Archive and Version Service.

On the stopping the PostgreSQL instance of Archive service, the database status still displays running.

Explanation: The PostgreSQL instance of Archive service is stopped. When you run `rcpostgresql status` PostgreSQL status may display as running. This is because the `rcpostgresql` script checks for any running instance of PostgreSQL database.

The database port change is not reflected on modifying arkConfig.xml file.

Possible cause: If you have configured Archive & Version Service using YaST utility, the scripts are automated and linked to each other. To reflect changes to port number or username through manual edits, you have to modify both the `pg_restart.sh` and `arkConfig.xml` file.


Action: For example, to change the port number to 2456, do the following:

1. Edit the `arkConfig.xml` located at `/etc/opt/novell/arkmanager/conf/`
 - ♦ Replace the number between the `portNumber` tags
`<portNumber>2456</portNumber>`
2. Edit `pg_restart.sh` file located at `/opt/novell/arkmanager/bin`
 - ♦ Replace the port number after `-p` option in the following line:

```
su dbuser4 -c "postmaster -D /var/opt/novell/arkmanager/data -h 127.0.0.1 -p 2456 -i"
```

This changes the database to run on port 2456.

Browse icon not functional for restoring files with NetStorage in a Firefox browser

Explanation: When you are restoring file versions with NetStorage in a Firefox browser, the Browse icon  that allows you to browse for the location of the file to be restored is not functional.

Action: Enter the path manually in the *Restore to* field.

C.2 Migration Issues

- ♦ [“Migration fails with the error message “PostgreSQL authentication failed.”” on page 123](#)
- ♦ [“On migrating to OES 11 server, Archive server fails to start.” on page 123](#)
- ♦ [“Migrating Archive server to OES 11 fails randomly” on page 123](#)
- ♦ [“Migration fails with the error message, “Unable to open ArkConfig.xml file”.” on page 123](#)

Migration fails with the error message "PostgreSQL authentication failed."

Possible cause: The username, password or database port number is invalid.

Action: Ensure that the PostgreSQL credentials are valid.

On migrating to OES 11 server, Archive server fails to start.

Possible cause: Linux operating system is case sensitive. Archive server fails to start if the archive path is not in appropriate case.

Action: Ensure the volume name is in uppercase. For example, *ARKVOL:\archive_directory*

Migrating Archive server to OES 11 fails randomly

Possible cause: The migration tool fails randomly to migrate Archive and Version Services to OES 11 server.

Action: To migrate Archive and Version Services, run the `sys:\system\arkmig.ncf` file on your NetWare server, then run the migration tool on the OES 11 server. Archive and Version service migrates successfully.

Migration fails with the error message, “Unable to open ArkConfig.xml file”.

Possible cause 1: `ArkConfig.xml` file name is case sensitive. Migration fails if the file name is not appropriate.

Possible cause 2: The spelling of the `ArkConfig.xml` file is incorrect.

Action: The file name is case-sensitive. Ensure you have specified the file name as `ArkConfig.xml`.

C.3 NSS File Version Utility

- ♦ [“Restoring versions, using NSS File Version Utility on Windows workstation, displays an error, unable to locate MSVCR71.dll file.” on page 124](#)
- ♦ [“How to debug issues with NSS File Version Utility.” on page 124](#)
- ♦ [“On installing NSS File Version Utility, Restore Archived Versions option is not available using GNOME desktop.” on page 124](#)

Restoring versions, using NSS File Version Utility on Windows workstation, displays an error, unable to locate MSVCR71.dll file.

Possible cause: MSVCR71.dll is not available in c:\windows\system32 folder.

Action: On installing Novell Client, the MSVCR71.dll file is installed on your system. Copy this file to c:\windows\system32 location. You will now be able to restore versions.

How to debug issues with NSS File Version Utility.

Possible cause: By default, the log file includes only fatal errors.

Action: To log debug or all messages, you have to configure the NovellVersioningClient.exe.config file and edit the level value parameters to DEBUG or ALL.

```
<level value="DEBUG"> or <level value="ALL">
```

For more information on editing log files, refer the [Apache Logging Services website \(http://logging.apache.org/log4net/\)](http://logging.apache.org/log4net/)

- On Windows 2000 and Windows XP, the default location for the log.txt file is Program Files\Application folder. For example, C:\Program Files\Novell\Versioning\log.txt.
- On SLED 10 and SLED 10 SP 1, the default location for the log.txt file is \$home/.versioning folder.
- In Windows Vista 32-bit platform the log file is located at c:\users\<username>\AppData\Local\VirtualStore\Program Files\Novell\Versioning.
- In Windows Vista 64-bit platform, to create log file, you need to configure the c:\Program Files\Novell\Versioning \NovellVersioningClient.exe.config file.

On installing NSS File Version Utility, Restore Archived Versions option is not available using GNOME desktop.

Action: When using GNOME desktop, you must login again to the Nautilus file manager to access the Restore Archived Versions option.

On un-installing NSS File Version Utility, Restore Archived Versions option is still displayed in the menu using GNOME desktop.

Action: When un-installing NSS File Version Utility from GNOME desktop, you must login again to the Nautilus file manager to clear the Restore Archived Versions option from the menu list.

The warning window in KDE terminates, if you click the button *Keep Running*.

Action: On KDE, if you try close the NSS File Version Utility window when in a processing state, a warning is displayed.



If you click either the Terminate or Keep Running button, the process is terminated. This behavior occurs because KDE kills the process, assuming it is not responding.

