Solstice Backup™ 6.1

Administrator's Guide

Please Recycle

Adobe PostScript

Contents

Figures

Tables

Preface

The *Solstice Backup 6.1 Administrator's Guide* contains information on how to configure and manage the Solstice Backup™ software.

This guide also provides information about how to use and administer the Solstice Backup software when it is installed in a typical networked environment. For more detailed technical information about Solstice Backup commands, refer to the online Solstice Backup manual (man) pages after you install the software.

The information in this guide is intended for system administrators who are responsible for installing software and maintaining the servers and clients on a network. Operators who monitor the daily backups may also find this manual useful.

Before You Read This Book

In order to fully use the information in this document, you must have thorough knowledge of the topics discussed in these books:

- *Solstice Backup 6.1 Installation Guide*
- *Solstice Backup 6.1 Release Notes*

Using UNIX Commands

This document may not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following for this information:

- *Solaris Handbook for Sun Peripherals*
- AnswerBook2™ online documentation for the Solaris™ operating environment
- Other software documentation that you received with your system

Typographic Conventions

TABLE P-1

Typeface	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your .login file. Use ls -a to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% **su** Password:
AaBbCc123	Book titles, new words or terms, words to be emphasized	Read Chapter 6 in the *User's Guide*. These are called *class* options. You *must* be superuser to do this.
	Command-line variable; replace with a real name or value	To delete a file, type rm *filename*.

Shell Prompts

TABLE P-2

Shell	Prompt
C shell	*machine_name*%
C shell superuser	*machine_name*#
Bourne shell and Korn shell	$
Bourne shell and Korn shell superuser	#

Related Documentation

TABLE P-3

Application	Title	Part Number
Installation	*Solstice Backup 6.1 Installation Guide*	816-1413-10
Release Notes	*Solstice Backup 6.1 Release Notes*	816-1414-10

Accessing Sun Documentation Online

The `docs.sun.com`^{SM} web site enables you to access a select group of Sun technical documentation on the Web. You can browse the `docs.sun.com` archive or search for a specific book title or subject at:

`http://docs.sun.com`

Ordering Sun Documentation

Fatbrain.com, an Internet professional bookstore, stocks select product documentation from Sun Microsystems, Inc.

For a list of documents and how to order them, visit the Sun Documentation Center on Fatbrain.com at:

`http://www.fatbrain.com/documentation/sun`

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can email your comments to Sun at:

```
docfeedback@sun.com
```

Please include the part number (816-1412-10) of your document in the subject line of your email.

Introduction

The Sun Solstice Backup client/server environment provides the ability to protect your enterprise against the loss of valuable data. In a network environment, where the amount of data grows rapidly as computers are added to the network, the need to protect data becomes crucial. The Sun Solstice Backup product gives you the power and flexibility to meet such a challenge.

This chapter introduces the Solstice Backup software and provides background information, such as terms and key concepts, required to use and understand the product. This chapter also includes a summary of the printed and electronic documentation available to you. Subsequent chapters in this guide provide specific information on Solstice Backup interfaces, features, and methods for saving and recovering your data.

This chapter does not contain installation information. For instructions about installing the software, refer to the *Sun Solstice Backup Installation Guide*. The information contained in this *Administrator's Guide* assumes that the Solstice Backup software has already been successfully installed on your Solstice Backup server and clients.

The main sections in this chapter are:

- "Solstice Backup Storage Management" on page 44
- "How the Solstice Backup Software Works" on page 47
- "How the Solstice Backup Software Backs Up Data" on page 53
- "How the Solstice Backup Software Recovers Data" on page 55
- "Solstice Backup Product Family" on page 58
- "Solstice Backup Documentation" on page 60

Solstice Backup Storage Management

The Solstice Backup product is a network storage management application that backs up and recovers data across an entire network of computers. The software protects your organization's valuable data by automatically backing up, storing, and indexing each individual data item (disk partition, directory, or file), allowing you to recover any particular data item, or even the complete contents of a client. You also have the option of recovering older versions of existing data items.

As the network and amount of data grow, Solstice Backup software provides the capacity and performance to handle the load. It offers a variety of features that give you the flexibility to sort your backed-up data by groups of specific clients, or onto specific storage volumes. It also provides regular backup status information regarding your data, media, and devices. Because the Solstice Backup index files that track the backed-up data and media can grow large over time, old and outdated entries are removed, based on your criteria, to keep the index size manageable.

Solstice Backup software features include a storage management application that directs high performance to a wide range of storage devices. The Solstice Backup software product is made up of the following components:

- Solstice Backup client software, which is a collection of processes and programs installed on the computers that contain data to be managed.

- Solstice Backup server software, which is a collection of processes and programs installed on the computer that performs the Solstice Backup services, including data backup, recovery, and archiving.

- Solstice Backup storage node software, which is a collection of processes and programs installed on computers that control their attached storage devices during Solstice Backup operations, including backup, recovery, and archive.

You can configure your Solstice Backup software to back up data to storage devices that are directly connected to the Solstice Backup server, or you can create a Solstice Backup *storage node*, which is a separate computer with one or more storage devices connected. The Solstice Backup server maintains the resources and files that keep track of Solstice Backup data. The storage node computer controls the storage devices and media.

You can direct backup, archive, and migration data to specific collections of volumes, according to particular data characteristics. As data is written to storage media, the Solstice Backup server creates a database to track the location of all managed data. The Solstice Backup server monitors data status and automatically rotates data through the tracking system as it reaches different levels of obsolescence, as defined by you.

You and your users can browse an index of recoverable files, create reports that describe the status of data or media, and recover data to a point in time specified by the user. In response to a data recovery request, the Solstice Backup storage management system locates the volume that contains the requested data and either directs a device to mount the appropriate volume for recovery, or sends a message to the operator to request the volume by name. In this way, the Solstice Backup system manages all storage volume operations.

Solstice Backup software is managed and operated using a combination of the following three methods:

- By using the Solstice Backup Administrator graphical interface to manage and configure scheduled backups, add clients and devices, and configure the Solstice Backup server. Using the Solstice Backup Administrator program, you can configure and perform backup and recovery tasks for any Solstice Backup server, client, or storage device on the network.

- By using the Solstice Backup Backup, Restore, Archive, and Retrieve graphical interfaces to manually backup, recover, and archive data between your client computer and your Solstice Backup server. Solstice Backup servers need to have these programs installed along with the Administrator program because the server acts as a client of itself.

- By running Solstice Backup command-line utilities. The software requires that certain tasks, such as autochanger and silo device configuration, be done with command-line utilities.

Note – The term "autochanger" refers to a variety of backup devices: autoloader, carousel, datawheel, jukebox, library, and near-line storage.

Cross-Platform Management

Solstice Backup software is a cross-platform storage management product. You can administer a Solstice Backup server from a computer running Microsoft Windows, UNIX, or Novell NetWare if the computer has the Solstice Backup client software installed on it and is connected to a network. Additionally, a Solstice Backup client on one platform can back up its data to a Solstice Backup server on a different platform.

You can direct and conduct administration services for any client or server on the network that has the Solstice Backup software installed. By displaying only the options that are relevant for that computer, the software frees you from the need to know the configuration of individual computers before you start a monitoring session.

Performance

Standard Solstice Backup performance features include:

- Server parallelism, which allows several *save streams* to flow to the server or storage node at the same time.
- Multiplexing, which allows more than one save stream to write to the same device at the same time.
- Client parallelism, which allows the client to send more than one save stream at a time.
- Session management, which allows you to control the number of save streams per device to maximize the performance of each device.
- Backup service to file devices and optional subsequent staging to near-line or offline volumes.

Optional Solstice Backup Software Additions

Optional Solstice Backup software additions include:

- Networker Autochanger Module
- Networker Silo Software Module
- Networker Archive Module
- Networker Database Modules (for backing up several types of databases)
- Legato Networker SNMP (Simple Network Management Protocol)
- Legato GEMS SmartMedia

Ease of Use

Solstice Backup software provides tools to make protection of critical data easy to manage. With these tools, you can:

- Use either the graphical interfaces or command-line programs to manage Solstice Backup tasks and functions.
- Administer and configure Solstice Backup functions from any computer on the network.
- Grant permission to provide the capability for recovery of one client's data to another client computer.
- Obtain immediate answers to questions by accessing online help and man pages.

- Take advantage of the automatic media management feature to allow the Solstice Backup server or storage node to label and mount volumes as needed for backups.
- Use the Tech Dialog and technical bulletins on the Sun web site to find answers to common questions.

Scalability

Your current Solstice Backup software can be scaled as your storage management needs grow. For example, you can:

- Upgrade the basic level of server functionality, add support for additional (or larger) autochangers, add support for more clients, or add optional software modules without the need to reinstall the server software.
- Add special Legato NetWorker Module client software to back up databases and other non-filesystem data.
- Add support for remote storage nodes to control backup devices, while the data management tasks remain centralized on a controlling Solstice Backup server.

How the Solstice Backup Software Works

Solstice Backup software works as a client/server system consisting of Solstice Backup servers, clients, storage nodes, storage devices, and storage volumes. The Solstice Backup *server* is a computer that receives backup data from Solstice Backup *clients*, indexes each data item, and sends the data to a storage device. Data from a single client computer backed up onto storage media is known as a *save set*. A save set typically comprises all of the backed-up data from a single filesystem or local disk. When the Solstice Backup server receives one or more client save sets, the data is written to a storage volume (usually a tape). The Solstice Backup server also indexes the contents of each individual volume to facilitate future recovery.

The Solstice Backup *client* is any computer that has the Solstice Backup client software installed, and that backs up its save sets to a Solstice Backup server. A Solstice Backup client can be a workstation or a dedicated server. It can also have a different platform or operating system from the Solstice Backup server used for backing up its data. A Solstice Backup server is always a client of itself, and Solstice Backup servers can also back up to other Solstice Backup servers as part of a hierarchical backup structure.

The index of backed-up data items is called the *client file index*. The index of volumes is called the *media database*. Collectively, these indexes are called the *online indexes*, and they form the backbone of the Solstice Backup software's ability to locate and recover data. When you submit a request to view or browse your backup data, the Solstice Backup server looks to the client file index for a record of each backed up data item. In cases where the data item has been backed up several times over a period of time, the client file index provides this information, and gives you the option of recovering different versions of the same data item.

When you request recovery of a selected data item, the Solstice Backup server searches the media database to locate the specific volume or volumes containing the data items you want. If your storage device is a *stand-alone device* (a single-drive device that requires volumes to be manually inserted), the Solstice Backup server requests you to insert the appropriate volume into the device. If your storage device is an *autochanger*, the server instructs the autochanger to locate the volume in the appropriate autochanger slot.

Once you select the data you want to recover and the Solstice Backup server locates the appropriate volume, the data is written back to the original client location; or in certain cases, such as *directed recovery*, to other client locations you choose. In cases where recovered data might overwrite data already on the client, the Solstice Backup server alerts you and gives you several options for solving the conflict. You can also configure in advance the action to take when recovery conflicts occur.

shows the Solstice Backup data flow, starting with the Solstice Backup client and ending with the storage volume.
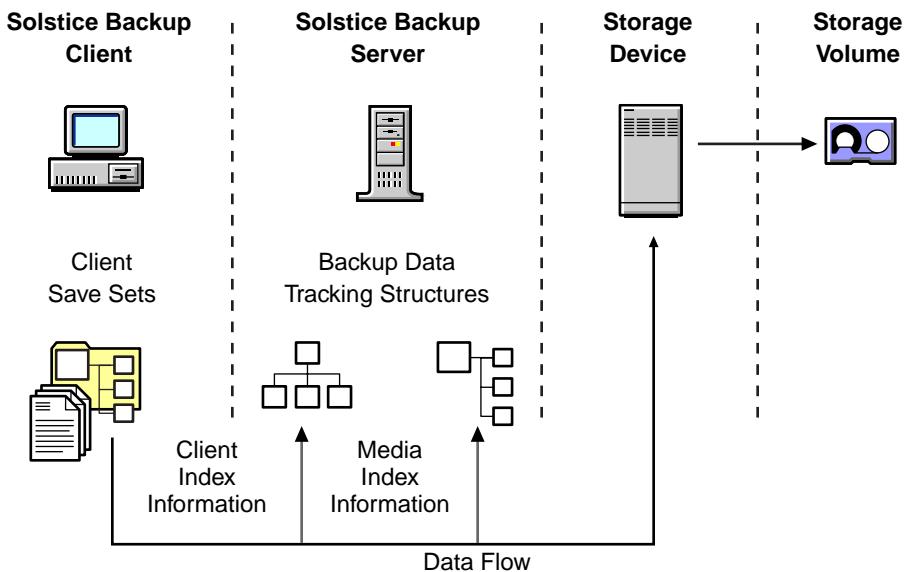
FIGURE 1-1 Data Backup Flow

Solstice Backup Resources and Attributes

Solstice Backup performance and functionality are configured through Solstice Backup *resources*. Resources are the building blocks used by the Solstice Backup server to license, configure, and operate Solstice Backup clients and servers to back up, store, and recover data. Solstice Backup resources are represented in the Administrator program using various windows.

Most resources can be created, edited, and deleted by administrators with the proper permissions. The dialog boxes for these resources contain the attributes for creating and editing resources. Solstice Backup *attributes* are the individual components of a resource. Attributes are represented by the various text boxes, drop-down lists, and so on that are contained in a resource dialog box. For example, a client would be a Solstice Backup resource, and the client's name would be an attribute of that resource.

The Solstice Backup product comes packaged with preconfigured and default resource settings. However, you can modify these resources to optimize the performance of your server. Many of the features and commands found in the Solstice Backup interface, and much of the information contained in this *Administrator's Guide*, pertain to resource creation and configuration.

lists Solstice Backup resources and their descriptions.

TABLE 1-1   Solstice Backup Resources

Resource	Action
Autochanger	Describes a device that can mount and write to multiple volumes automatically. See "Device" in this table. The term "autochanger" refers to a variety of backup devices: autoloader, carousel, datawheel, jukebox, library, and near-line storage.
Client	Identifies a Solstice Backup client, including data items such as filesystems, disk volumes, directories, or files that will be backed up, and assigns specific backup instructions to the data.
Device	Describes a device that is connected to the server or storage node.
Directive	Defines file processing activities that are performed during a backup of a client.
Group	Names and defines a unique collection of Client resources that participate in a group backup operation. Clients can follow different backup schedules within the group operation, but all clients in the group start the operation at the same time.
Label Template	Provides instructions for the Solstice Backup server to follow whenever the storage management application needs to create a label for a volume. Each pool is associated with a specific label template.
Notification	Names and defines a specific Solstice Backup event and an associated response, such as a message logged to a log file or sent to an administrator via email.
Policy	Names and defines a time period that can be associated with administrator-specified data lifecycle policies.
Pool	Names a collection of volumes and specifies the acceptance criteria that backup data must meet to be written to this named collection.
Registration	Defines the features enabled for a particular Solstice Backup server.
Schedule	Names and defines a backup schedule that can be assigned to a Client or Group resource. The schedule dictates the days that the data is to be backed up and the level of each backup.
Server	Identifies the local Solstice Backup server characteristics. Only one Server resource exists per Solstice Backup server.
Staging	Names and defines the Solstice Backup policy for moving data from one storage medium to another, and for removing the data from its original location.

Many of these resources also exist as attributes of other resources. These interdependent relationships affect the order in which you implement the configuration. Certain component resources must be defined before they can be listed as attributes of other resources. For example, to apply a custom Label Template resource to a pool, you must create the label template before configuring the pool.

For more detailed information on accessing and using Solstice Backup resources and attributes in the Solstice Backup Administrator program, see "Solstice Backup Administrator Interfaces" on page 78.

Solstice Backup Daemons and Programs

Solstice Backup client/server technology uses a network Remote Procedure Call (RPC) protocol to back up your data. Solstice Backup software consists of several daemons and programs that oversee the backup and recover processes, as well as storage management client configurations, a client file index, and a media database.

The Solstice Backup environment calls upon several daemons and programs when a request for a backup is received. The daemons coordinate the tasks associated with a backup or recovery, and record information about what was backed up, as well as the media containing the backed-up data.

TABLE 1-2 on page 51 provides a description of the Solstice Backup server and storage node daemons and programs. The Solstice Backup man pages contain further details about the Solstice Backup daemons and programs.

TABLE 1-2 Server Daemons and Programs *(1 of 2)*

Daemons or Program	Function
ansrd	Monitors an active save or recover session; agent process spawned by nsrd in response to a save or recover session.
asavegrp	Monitors the progress of individual save sets; agent process invoked by the savegrp program.
nsrck	Checks the consistency of the client file index; invoked by nsrd whenever the consistency of the client file index needs to be confirmed.
nsrd	Provides an RPC-based save and recover service to Solstice Backup clients. This is the master Solstice Backup daemon. For details regarding this daemon, see "The Master Solstice Backup Server Process" on page 52.
nsrim	Manages the server's client file index; invoked at the end of an instance of the savegrp program.
nsrindexd	Provides a method for inserting entries in the client file index based on information passed by the save program.
nsrmmd	Provides device support, generates mount requests, and *multiplexes* save set data during a multiclient backup (nsrd can start several nsrmmd daemons, up to the number specified in the device's Target Sessions attribute). This is the media multiplexor daemon.

TABLE 1-2   Server Daemons and Programs *(2 of 2)*

Daemons or Program	Function
nsrmmdbd	Provides media and save set database management services to the local nsrd and nsrmmd daemon and records entries in the media database. This is the media management database daemon.
savegrp	Runs a group of Solstice Backup clients through the save process.

The Master Solstice Backup Server Process

The nsrd process is the master Solstice Backup server daemon, responsible for the following tasks:

- Starting other daemons
- Allocating media daemons on server and storage node machines
- Authorizing backup and recover services for the client
- Contacting clients for scheduled backups
- Maintaining Solstice Backup configuration information
- Monitoring backup and recover sessions
- Maintaining server statistics and message logs

The nsrd Solstice Backup server daemon calls on the nsrexecd Solstice Backup client daemon and several client-side programs when it receives a scheduled or manual backup request. The ansrd agent starts on the Solstice Backup server computer to monitor the progress of the backup session.

TABLE 1-3 describes the client-side daemon and programs.

TABLE 1-3   Client Daemons and Programs

Daemons or Program	Function
nsrexecd	Authenticates the Solstice Backup server's remote execution request and executes the save and savefs programs on the client.
recover	Browses the Solstice Backup server's client file index and restores the specified file to primary disk storage.
save	Sends specified files in a multiplexed data stream to the Solstice Backup server for backup to media by nsrmmd and entry in the client file indexes and media database by nsrindexd.
savefs	Saves a filesystem (using the save program) to a Solstice Backup server; a level-based system is used to save only those files which have been modified since some previous save (a partial save).

How the Solstice Backup Software Backs Up Data

When you configure a *backup group* on the Solstice Backup server, you schedule a start time for the backup group. FIGURE 1-2 on page 53 shows how all the Solstice Backup client and server daemons and programs interact during a scheduled backup.
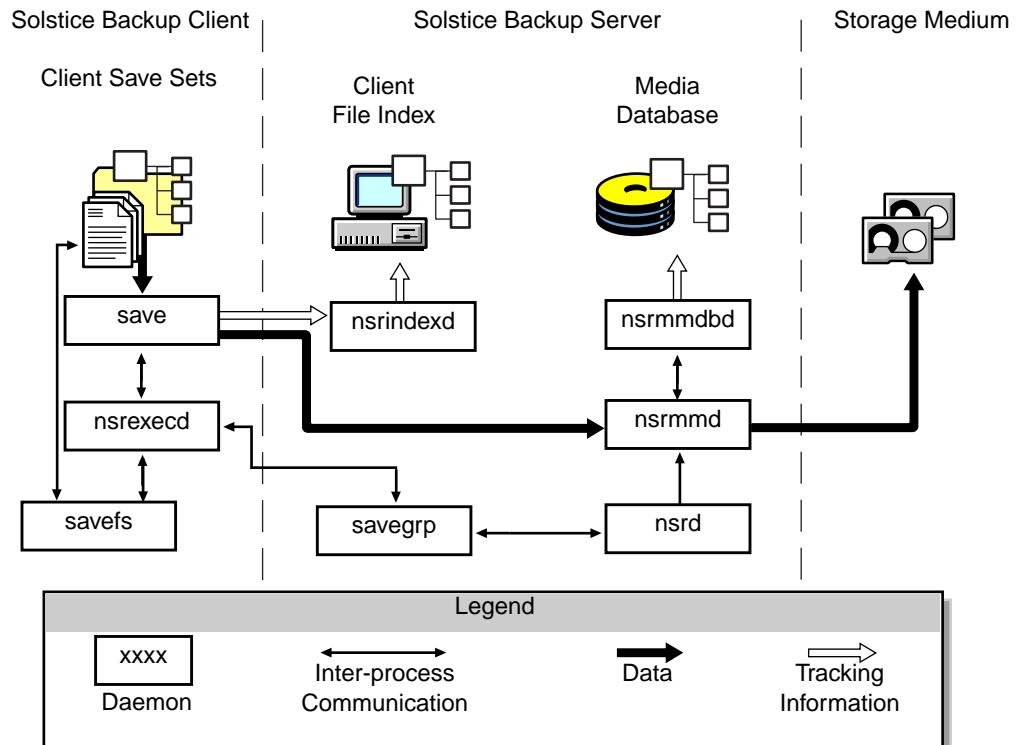
FIGURE 1-2   Daemons and Programs during a Scheduled Backup

The nsrd server daemon starts the server's savegrp program for the backup group at the scheduled time.

The savegrp program queries all Client resources configured on the Solstice Backup server to determine the following:

■ Which clients configured on the server are members of the scheduled group

- What level of backup (`save`) to perform
- How many save sets to run concurrently, determined by the parallelism value set on the Solstice Backup server
- When the most recent backup of the group occurred

If any of this information is not available on the Solstice Backup server, `savegrp` sends a request to the `nsrexecd` client-side daemon to run `savefs` on each client assigned to the backup group to gather the necessary details.

The `savefs` program tells `savegrp` which objects to back up for the client. After `savegrp` receives information about the objects to back up, `savegrp` assembles a work list for the server. The work list specifies the order in which clients are contacted for backup. The order of the work list is determined by the Client Priority attribute in the Client resource. The client with the lowest value in the Client Priority attribute is contacted first.

If problems were encountered with the client file index during the previous backup session, `nsrd` invokes the `nsrck` program to check the consistency and state of the Solstice Backup server's client file indexes and media database. Then `nsrd` starts the `nsrindexd` client file index insertion daemon.

The `savegrp` program contacts the first client on the server's work list. The client's `nsrexecd` is invoked and starts a `save` session of the first save set listed on the server's work list. The `save` program passes to `nsrd` all save criteria, such as group, client, save sets, storage node affinity, and level of the save data. With this information, `nsrd` determines the pool of volumes that will store the data and forwards the information to the appropriate media daemon, on either the Solstice Backup server or a storage node controlled by the Solstice Backup server.

As a result, the `nsrmmd` media daemon:

- Sends a message to the console of the Solstice Backup server, requesting a mount of the media assigned to the volume pool indicated by `nsrd`.
- Writes the data sent by `save` to storage media.
- Forwards storage information to `nsrmmdbd` for recording in the Solstice Backup server's media database.

Any time there is a lull in save set activity from the client, the Solstice Backup server attempts to find another save set in the group to keep the process moving. The `savegrp` program attempts to concurrently back up as many save sets as possible, up to the limit set by the Parallelism attribute in the Solstice Backup server's configuration.

The `savegrp` program repeats the process for each item on the server's work list until all clients in the group are backed up. If the Solstice Backup server is part of the group being backed up, `nsrim` is invoked before the `savegrp` is completed and the Solstice Backup server's *bootstrap* file is backed up. The bootstrap file contains essential information for the Solstice Backup disaster recovery procedures. When the

bootstrap backup is completed, a bootstrap printout is sent to the default printer configured for the Solstice Backup server. Keep the bootstrap printout in a safe place in case you need to restore the Solstice Backup server.

The final results of the savegrp execution are sent back to the server and are included in the savegroup completion report, which is sent by e-mail to the root user.

How the Solstice Backup Software Recovers Data

FIGURE 1-3 shows how the Solstice Backup server and client daemons and programs interact while recovering data to a Solstice Backup client.
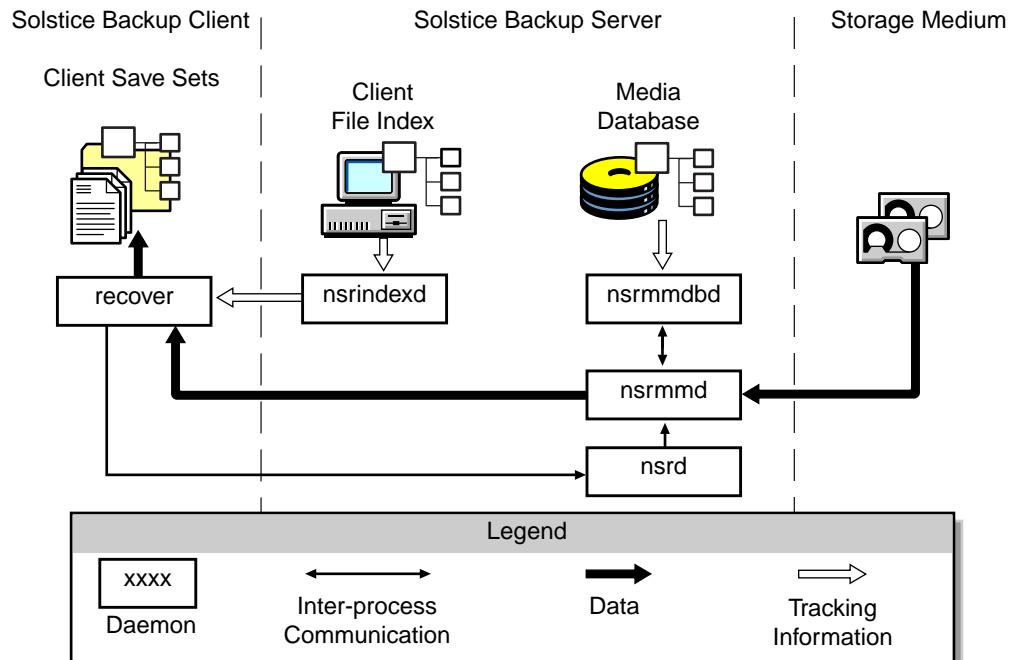
FIGURE 1-3 Daemons and Programs during a Recover Session

When the Solstice Backup server receives a `recover` request from a client, the server's `nsrd` daemon contacts the server's `nsrmmd` media daemon. The `nsrmmd` daemon contacts the server's `nsrmmdbd` media database daemon to determine which media contains the save set requested by `recover`. After the save set's media location is obtained, `nsrmmd` issues a mount request, the media is positioned to the beginning of the save set, and the save set stored on the mounted media is passed to `nsrmmd`. The media daemon forwards the save set to the client's `recover` program, which restores the data to the client's filesystem.

When the server's `nsrmmdbd` media database daemon cannot locate the required volumes, or when there are no enabled devices on the Solstice Backup storage node or server, the `recover` request fails. The Solstice Backup software displays the following message on the client machine that requested the `recover`:

```
NSR server client: no matching devices on server or storage node
```

The Solstice Backup software also displays the following message for the Solstice Backup server:

```
media notice: no matching devices on server or storage node for recover
by client client_name

media notice: enable or check device(s) on server or storage node
```

Solstice Backup Operations with Storage Nodes

A *storage node* is a machine that contains Solstice Backup client and media management daemons and programs. A storage node is connected to one or more storage devices that are used in Solstice Backup operations, such as backup and recovery. Media management daemons on the storage node machine read and write data on storage volumes.

Backup and recovery operations with storage nodes are very similar to backup and recovery operations on the Solstice Backup server. The major difference is where the data resides.

A storage node is controlled by a Solstice Backup server (or multiple Solstice Backup servers). The Solstice Backup server's nsrd master daemon starts an instance of nsrmmd (media management daemon) on the storage node. The Solstice Backup server routes appropriate data to each media management daemon according to the Storage Node attribute in the Client resource.

After data is written to the storage devices connected to a storage node, the Solstice Backup server tracks the location of the data. On the Solstice Backup server, the nsrindexd daemon writes entries in the client file indexes to track the location of each file in a save set, and the nsrmmdbd daemon writes entries in the media database to track the location of each save set on the media.

FIGURE 1-4 on page 57 shows the client's save process contacting the storage node's nsrmmd process.
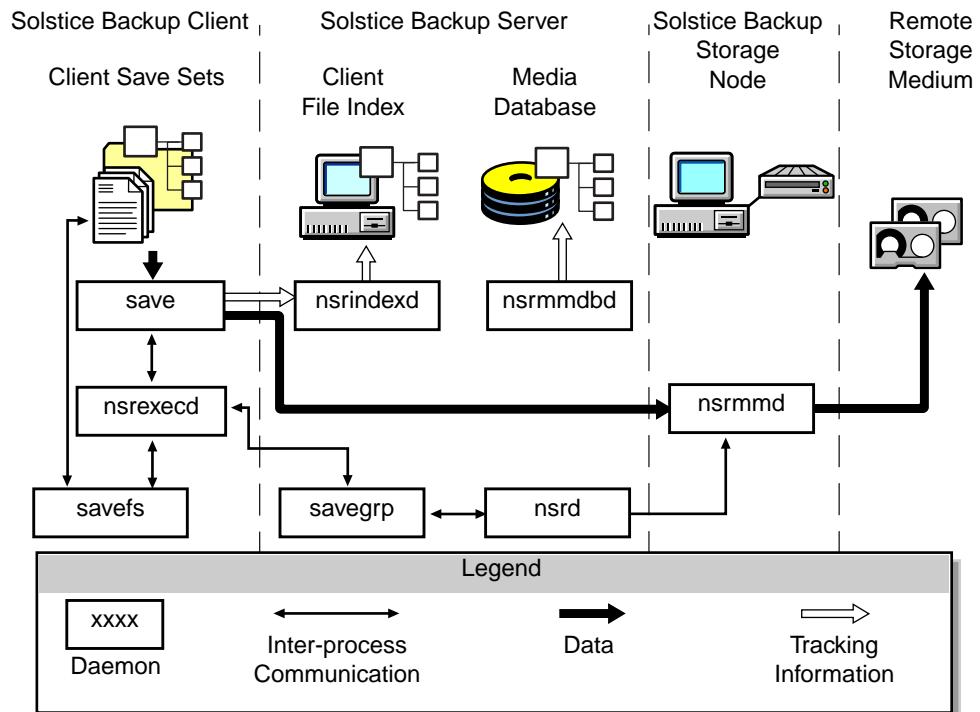
FIGURE 1-4   Daemons and Programs during a Save Session with a Storage Node

Recovery Using Storage Nodes

When the Solstice Backup server receives a request from the client's `recover` program, it contacts the `nsrmmd` daemon on the storage node. The daemon reads the data from media and sends the data to the Solstice Backup client, where the `recover` program rebuilds the data.

shows an example of how a `recover` process works in a Solstice Backup setup that includes a storage node.
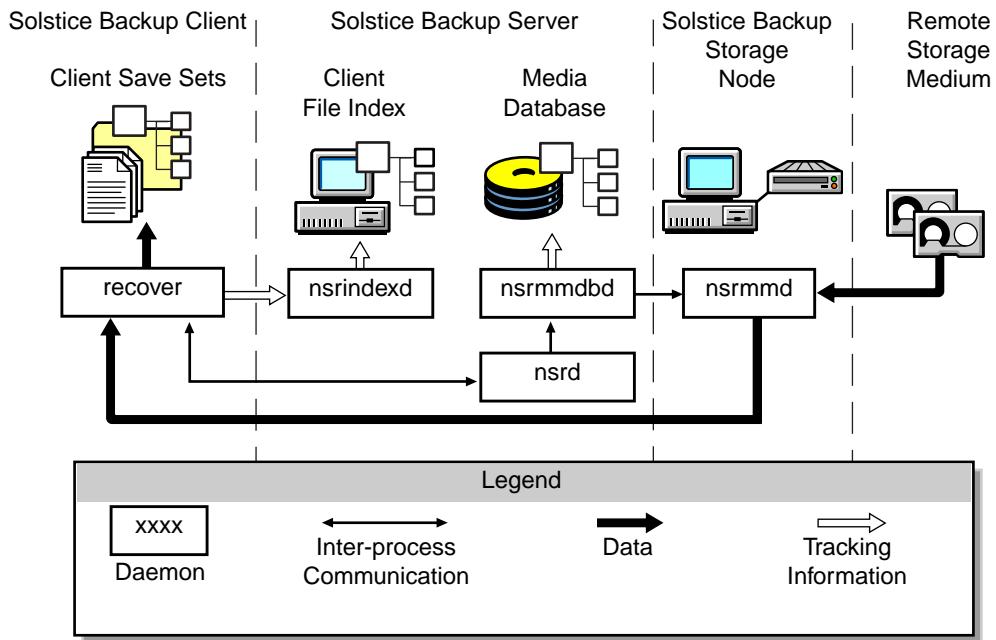
FIGURE 1-5   Daemons and Programs during a Recover Session with a Storage Node

Solstice Backup Product Family

Solstice Backup has a large product family of server, client, database, and related backup-and-recover software.

Server Software Products

All Solstice Backup clients, as well as Solstice Backup servers—which can function as their own clients and back up their own data—can be backed up by Solstice Backup server software.

Solstice Backup server products are available in three distinct versions offering different levels of support:

■ Solstice Backup WorkGroup Edition enables you to back up the server and up to three clients. Solstice Backup WorkGroup Edition does not include support for additional clients or optional software features, such as Networker Archive. If you decide at a later date to purchase a higher level of Solstice Backup software, all of the data backed up by Solstice Backup WorkGroup Edition is recoverable by any level of Solstice Backup software you purchase.

■ Solstice Backup Network Edition allows you to back up the server and as many clients as the number of client connections you have purchased. You can upgrade Solstice Backup Network Edition to include support for additional clients and autochangers, as well as optional software features.

■ Solstice Backup Power Edition is a server software product that has been specially tuned to protect and provide backup for environments that support Very Large Database (VLDB) or large filesystem applications (in the terabyte range). You can also upgrade Solstice Backup Power Edition to include support for high-speed devices.

Client Software Products

Some Solstice Backup client software products are included with the server products. Refer to your *Installation Guide* for any clients packaged with your particular server software and instructions on how to install them. Support for platform-specific clients is also sold separately as Legato NetWorker ClientPak software. Sun supports clients for the following platforms and operating systems:

■ UNIX: Solaris, SunOS, HP-UX, AIX, Linux, Compaq Tru64, DYNIX/ptx, IRIX, and UnixWare

■ PC Desktops: Microsoft Windows NT Workstation, Microsoft Windows 2000 Workstation, Microsoft Windows 98, and Microsoft Windows 95

■ Microsoft Windows NT Server and Microsoft Windows 2000 Server

■ Novell NetWare/IntraNetWare

For complete information on Solstice Backup client software, refer to the documentation in your Solstice Backup software media kit.

Other Products

Solstice Backup offers other products to enhance your storage management needs:

TABLE 1-4   Additional Solstice Backup Products

Product	Description	For More Information
NetWorker Autochanger Module	Enables the Solstice Backup software to connect to one or more autochangers.	Refer to Chapter 10.
NetWorker Silo Software Module	Enables the Solstice Backup software to connect to one or more silos.	Refer to Chapter 10.
NetWorker Archive feature	Provides a sophisticated method to archive valuable data for long-term storage.	Refer to Chapter 8.
Legato NetWorker SNMP (Simple Network Management Protocol) Module	Provides communication of Solstice Backup event notifications to SNMP-compliant network management stations.	Refer to Chapter 12.
Legato NetWorker Modules	Enables reliable backup and restore capabilities for the most widely used databases, including Microsoft Exchange Server and SQL Server, Oracle, INFORMIX, Lotus Notes, SAP R/3, and Sybase.	Refer to the separate documentation for each module.

Sun continues to develop enhancements to the Solstice Backup product line, including additional server, client, database, and add-on products. For the latest product information, refer to the Sun web site (*www.sun.com*).

Solstice Backup Documentation

The Solstice Backup product comes with a complete suite of documentation to provide the information you need to optimally run the Solstice Backup software. The following sections briefly describe the Solstice Backup documentation products.

Note – Solstice Backup documentation is periodically updated between releases, and the most recent versions are posted on the Legato web site.

Documentation Suite CD-ROM

All Solstice Backup product manuals and guides are provided in PDF format on a separate Documentation Suite CD-ROM as part of the product Media Kit.

Printed Solstice Backup and Legato NetWorker Product Manuals

You can also obtain a Solstice Backup Documentation Kit, including printed copies of the Solstice Backup manuals, for a nominal charge. Contact your Solstice Backup sales representative for information on ordering Documentation Kits for a particular Sun product.

Administrator's Guide

This *Administrator's Guide* is your main resource for understanding how the Solstice Backup product works. It provides information on a wide variety of Solstice Backup topics, including:

- How the Solstice Backup software saves and recovers data.
- Configuring Solstice Backup servers and clients.
- Understanding and using Solstice Backup resources.
- Configuring backup devices and media, including autochangers and silos.
- Using specialized storage management tools and protocols, including:
  - Archiving
  - Cloning
  - Staging
  - Network Data Management Protocol (NDMP)
  - Simple Network Management Protocol (SNMP)
- Troubleshooting

Installation Guide

The *Installation Guide* provides the following information:

- How to install the software components on your server, client, and designated storage node computers
- How to configure your backup devices

- How to update or upgrade your current Solstice Backup software
- How to remove the software components

Review the *Installation Guide* before you install the Solstice Backup software. After the installation has completed, keep a printed copy of the *Installation Guide* in a safe place, so that it is available if you need to reinstall or remove the software later.

Note – The *Installation Guides* for the Legato NetWorker ClientPak for NetWare, UNIX, and Windows client software provide instructions for installing the client software on other operating systems that are not supported as Solstice Backup servers, but are supported as Solstice Backup clients.

Release Supplement

The *Release Supplement* provides late-breaking information on software bugs, workarounds, and other items not found in the other documentation sources. Depending on the Solstice Backup release, the *Release Supplement* may also contain information on new and improved features that were not included in this *Administrator's Guide*, the *Disaster Recovery Guide*, or other documentation.

Updated versions of the *Release Supplement* are available in the Document Library on the Legato web site (www.legato.com). Refer to the web site periodically to view the latest version of this and all other Legato documentation. Compare the dates printed on each version of the Release Supplement to ensure you have the latest one.

Online Help

The Solstice Backup Administrator program and other graphical user interfaces include an online help component. The online help contains specific information about the GUI, including context-sensitive help in each window that describes the purpose of the window and what information to enter in each field in the window. There are also help topics that describe Solstice Backup features as well as topics that provide step-by-step procedures for doing Solstice Backup tasks with the GUI.

To access the online help, select the Help menu displayed in the graphical user interface. You have four options:

TABLE 1-5   Online Help Menu Options

Option	Description
On Window	Provides information for attributes in the current window.
On Topic	Provides a list of help topics related to Solstice Backup software.
On Help	Provides information about how to use the online help.
On Version	Provides information on the release version number of the Solstice Backup software that is installed on your system.

The Solstice Backup Administrator also provides Help buttons in most of the dialog boxes. To display the help topic associated with a particular Solstice Backup feature, click the Help button in a dialog related to the feature.

Man Pages

The man pages are a reference for the syntax and function of the Solstice Backup commands that you issue from the command line. To view man pages, make sure that the MANPATH environment variable includes the path where your Solstice Backup man pages are installed. By default, Solstice Backup man pages are installed in /usr/share/man.

Disaster Recovery Guide

The *Legato NetWorker Disaster Recovery Guide* is a multi-platform guide that provides essential information about how to use Solstice Backup products to recover data from your Solstice Backup servers in case of a disaster. This guide also contains tips and information about preventive measures to safeguard data and prepare your backup policies to guard against a disaster before one strikes.

Review the *Disaster Recovery Guide* when you first install and configure the Solstice Backup software so that you are prepared with strategies to protect your data and operations. Keep a printed copy of the guide with your other disaster recovery information, so that it is immediately available in case of an emergency.

Performance Tuning Guide

The *Legato NetWorker Power Edition Performance Tuning Guide* discusses the enhancements available with Solstice Backup Power Edition and tells you how to configure and measure the performance of your Solstice Backup Power Edition server.

Error Message Guide

The *Legato Error Message Guide* describes messages and troubleshooting procedures for problems that you might encounter while using the Solstice Backup software.

Legato Compatibility Guide

The *Legato Compatibility Guide* provides the latest information on software and devices compatible with Solstice Backup software. This document is available on the Legato web site at `www.legato.com`

Command Reference Guide

A complete set of the Legato online man pages is available to reference the syntax and function of the Solstice Backup commands you issue from the command line.

The *Legato Command Reference Guide* describes the commands and resources that you might encounter while running Legato products. It is your best source of information on commands you enter at the command prompt, including command options.

For instructions on how to use the guide, including tips and troubleshooting, view the README file in the `DOCS\cmndref` directory on the Documentation Suite CD-ROM.

Getting Started

This chapter explains how to begin using the Solstice Backup software to perform storage management activities.

The main sections in this chapter are:

- "Preparing to Use the Solstice Backup Software" on page 65
- "Operating Solstice Backup Software for the First Time" on page 66
- "Verifying Your Operations" on page 72
- "Customizing and Configuring the Solstice Backup Software" on page 75
- "Solstice Backup Administrator Interfaces" on page 78
- "Solstice Backup Manual Backup, Recovery, and Archiving Interfaces" on page 84
- "Rehosting a Solstice Backup Server" on page 90

Preparing to Use the Solstice Backup Software

If you are new to using Solstice Backup software, you should go through the following steps:

1. Review the Solstice Backup documentation, particularly the *Installation Guide* and the most recent *Release Supplement*.

2. Install the Solstice Backup server and client software on your computers.

3. Perform the tasks listed in "Operating Solstice Backup Software for the First Time" on page 66. These tasks include:

   - Starting the Solstice Backup Administrator program.
   - Labeling and mounting a volume.

- Backing up your data.
- Adding clients.

Once you have completed these tasks, learn how to monitor and verify data during a backup as described in "Verifying Your Operations" on page 72, then continue to customize and configure the core Solstice Backup functions.

- Learn when and how to perform procedures for the following tasks described in Chapter 7:
  – Manual backup
  – Data recovery
  – Directed recovery
  – Save set-level recovery
  – Data recovery using the scanner program
- Review the discussion of Solstice Backup browse and retention policies and online indexes (client file index and media database). See "Browse and Retention Policies" on page 156.
- Review disaster recovery procedures. Refer to "Disaster Recovery" on page 299 and the *Legato NetWorker Disaster Recovery Guide*.

Reviewing and working with this material helps you develop an understanding of how Solstice Backup software works, and of the basic functions and features that you may customize later to optimize performance and reliability.

Operating Solstice Backup Software for the First Time

This section provides an overview of the software you use to operate the Solstice Backup software. It also presents a series of steps to run basic Solstice Backup functions. Whenever possible, the tasks in this section use default and preconfigured settings. In subsequent sections, as well as the remaining chapters of this *Administrator's Guide*, you will learn how to customize Solstice Backup resources and attributes, and how to configure the Solstice Backup software to work for your particular needs.

Operating Tools

To begin using Solstice Backup software, you need to first connect to a server, mount a volume for your backups, and enable a configuration. Then you can start a backup.

The primary graphical interface you use to operate Solstice Backup software is the Administrator program (nwadmin). You might also use any combination of the following:

- Solstice Backup backup program (nwbackup)
- Solstice Backup recover program (nwrecover)
- Solstice Backup archive program (nwarchive)
- Solstice Backup retrieval program, for retrieving archived data (nwretrieve)
- Command-line interfaces

The combination of these tools enables you to:

- Perform scheduled and manual backups
- Verify the servers you are using
- Verify that the necessary processes are running on their respective computers
- Verify the success of your backup
- Read message logs
- Perform recoveries

For a full description of these programs and their features, see "Solstice Backup Storage Management" on page 44 and "Solstice Backup Resources and Attributes" on page 49.

You should also become familiar with the directories listed in TABLE 2-1 (located in the Solstice Backup installation directory, usually /nsr):

TABLE 2-1   Solstice Backup Files and Directories

Directory	Files
mm	Solstice Backup media database
res	Solstice Backup resource files: • nsr.res • nsrjb.res • nsrla.res • servers
index/client_name/db	Client file index database
logs	Solstice Backup log files

As described in Chapter 1, you configure performance and functionality through Solstice Backup resources. If you do not want to create a customized configuration, you need only enable a group to begin backups according to the default configuration.

To learn how to customize your Solstice Backup resource settings, such as Group, Client, or Pool, see "To Customize a Backup" on page 70. For more detailed explanations, see "Customizing and Configuring the Solstice Backup Software" on page 75 and subsequent chapters in this guide.

▼ To Start the Solstice Backup Administrator Program

To start the Solstice Backup Administrator program and connect to your server, enter the following command at the prompt:

```
# nwadmin &
```

▼ To Label and Mount a Volume

To perform backups, your backup media must be configured and available for the Solstice Backup server to write data to it. Once you have started the Solstice Backup Administrator program, label and mount a blank media volume:

1. Insert a volume into the stand-alone device drive or slot 1 of the autochanger.

Note – Backing up data to an autochanger requires that you *first* install and enable the Autochanger Software Module.

---

2. From the Media menu, select Devices to open the Devices window.

3. Click Create.

4. Enter the name of the device, and select the media type.

5. Click Apply.

6. Click Label to open the Label or Autochanger Label dialog box.

7. Select a Pool.

   Selecting a pool automatically chooses a label for the volume.

8. If you are backing up to a stand-alone device, select Mount After Labeling in the Label dialog box.

9. If you back up to a device in an autochanger, enter 1 in both the First Slot and Last Slot attributes in the Autochanger Label dialog box.

10. Click Apply.

The Solstice Backup software labels the volume. If the storage device is a stand-alone device, the Solstice Backup software also mounts the volume and displays the name and status of the volume in the Monitor window. For example:

```
/dev/rmt/0cbn                         venus.001     labeling, done
```

Backing Up Your Data

Solstice Backup software contains preconfigured resources you can start using immediately; therefore, when the volume is successfully labeled, you are ready to perform a backup.

There are two types of backups in the Solstice Backup environment: scheduled and manual.

- Scheduled backups are configured in the Solstice Backup Administrator and are assigned a specific time for the backup to run, usually when network usage is low. For example, the preconfigured Default group is scheduled to begin its backup at 3:33 a.m. You can override scheduled backups to perform the backup immediately, if necessary.

- Manual backups are backups performed by the user of the client computer, using the nwbackup program.

To run a scheduled backup, you need only enable the default schedule.

▼ To Enable a Scheduled Backup

To configure the Solstice Backup server for a scheduled backup:

1. From the Customize menu, select Groups to open the Groups window.

2. In the Groups attribute, select the Default group.

3. For the Autostart attribute, select Enabled.

4. Click Apply.

Your Solstice Backup server is now enabled to back up itself, starting at 3:33 a.m. You can test the scheduled backup by overriding the scheduled time and starting the backup immediately.

▼ To Override a Scheduled Backup

To override a scheduled backup time and start the backup immediately:

1. From the Customize menu, select Groups to open the Groups window.

2. In the Groups attribute, select the Default group.

3. For the Autostart attribute, select Start Now.

4. Click Apply.

   Your Solstice Backup server immediately backs up itself.

Note – The scheduled backup will still occur at 3:33 a.m. But, by starting it now you are overriding the backup schedule so that you can see how the Solstice Backup server performs an automatic backup.

▼ To Customize a Backup

You can customize your backup schedule and other resources to suit your environment. Solstice Backup resources should be configured in a specific order, because some resources contain information that will appear in the Solstice Backup Administrator as attributes of other resources. Resources that are also attributes of other resources should be defined first. After this initial setup, you can then modify the resources at any point.

Because the Schedule and Policy resources do not depend on any other resources, customize these resources first. Then move on to the Group, Client, and Pool resources. FIGURE 2-1 on page 71 provides a graphical representation of the process for configuring Solstice Backup resources.
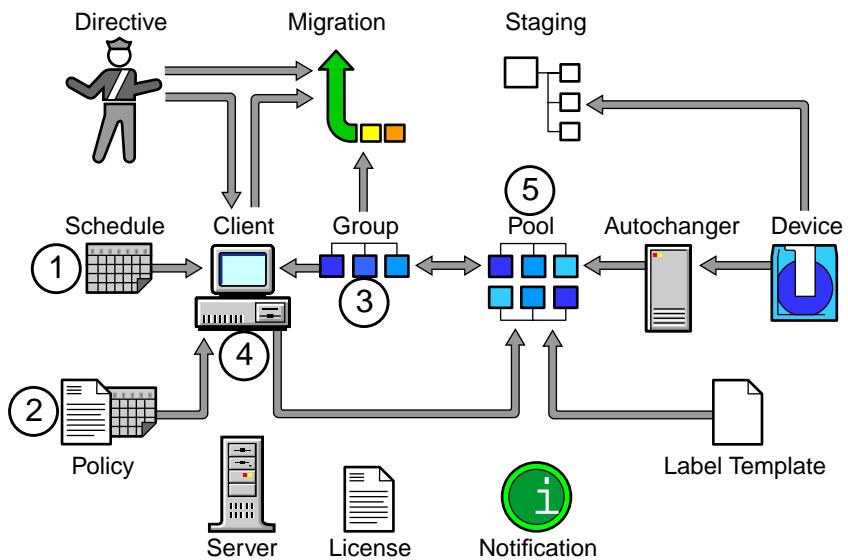
FIGURE 2-1   Resource Configuration Order

To customize backup resources:

1. Create your customized Schedule, which becomes an attribute of the Group resource. For details, see "Schedules" on page 268.

2. Create your customized Browse or Retention Policy, which becomes an attribute of the Client resource. For details, see "Browse and Retention Policies" on page 156.

3. Customize your Group resource, which becomes an attribute of the Client resource. For details, see "Backup Groups" on page 256.

4. Customize your Client resource, which can be an attribute of the Pool resource. For details, see "Client Configuration" on page 138.

5. Customize your Pool resource. For details, see "Pools" on page 194.

In FIGURE 2-1, resources with outgoing arrows are both independent resources and attributes of the resources to which the arrows point. Note that the Directive, Schedule, and Policy resources must be defined before they are available as attributes in the Client resource.

▼ To Add More Clients

Once you have successfully performed your first backup of the Solstice Backup server, you should add clients to the server. Because every Solstice Backup server is also a client to itself, you already have one client by default.

To add clients to the Solstice Backup server:

1. From the Clients menu, select Client Setup to open the Clients window.

2. Click Create.

3. In the Name text box, enter the new client's hostname.

4. Accept the default settings for the other attributes, or configure as appropriate.

5. Click Apply.

   When you successfully add a new client, the client's name and icon are displayed in the Clients window.

6. Repeat these steps for each Solstice Backup client you want to be backed up by your Solstice Backup server.

Verifying Your Operations

This section explains how to verify your backup. It also explains how to verify the authorized servers, how to verify that the Solstice Backup processes are running, and how to read message logs.

▼ To Monitor and Verify Data During a Backup

To monitor and verify the data during a backup:

1. From the Customize menu, select Groups to open the Groups window.

2. In the Groups attribute, select the Default group.

3. For the Autostart attribute, select Start now.

4. Click Apply.

In the Messages window of the main Solstice Backup Administrator window, the Solstice Backup server displays detailed messages about the progress of the backup.

▼ To View the Group Control Detail Dialog Box

You can also view information about the backup in the Group Control Detail dialog box.

To open the Group Control Detail dialog box from the main window of the Solstice Backup Administrator program:

1. Select Group Control.

2. Select Details.

The Solstice Backup Administrator program lists the successful, unsuccessful, and incomplete save sets in the Group Control Detail dialog box.

Note – If you have a stand-alone storage device, make sure that you leave the volume in the device; otherwise, the Solstice Backup software will not be able to begin the scheduled backup for the Default group at 3:33 a.m.

▼ To Verify Backup Servers

You can verify the servers that are authorized to back up a client computer by checking the client's /nsr/res/servers file, which lists the servers that are authorized to back up the computer.

Caution – If the /nsr/res/servers file contains no data (no servers are listed), any server can back up that computer, unless the client daemon (nsrexecd) is started using the -s *server* option.

▼ To Verify that the Solstice Backup Processes Are Running

To verify that the Solstice Backup processes are running:

1. Enter *one* of the following at the command prompt:

   - # **ps -ef | grep nsr**
   - # **ps -ax | grep nsr**

2. Depending on the type of installation, the processes listed in TABLE 2-2 should be running:

TABLE 2-2   Solstice Backup Processes

Type of Installation	Processes
Solstice Backup Server	nsrd - Solstice Backup server nsrexecd - Solstice Backup client (two processes should be running) nsrmmd - Solstice Backup media multiplexer (one process per device) nsrmmdbd - Solstice Backup media management database nsrindexd - Solstice Backup file index database process
Solstice Backup Client	nsrexecd - Solstice Backup client
Solstice Backup Storage Node	nsrexecd - Solstice Backup client nsrmmd - Solstice Backup media multiplexer (one process per device)

How to View a Message

Messages appear in the Solstice Backup Administration program and are also written to message logs.

To view a message:

■ Messages relating to the Solstice Backup server to which you are currently connected are displayed in the Messages window of the Solstice Backup Administration program. You can also use the command-line program, nsrwatch, to monitor these messages.

■ In message logs:

  – /nsr/logs/daemon.log (messages logged during a backup, and also displayed in the Sessions window of the Solstice Backup Administration program)

  – /nsr/logs/messages (messages logged after a backup is complete, and also displayed in the Messages window of the Solstice Backup Administration program)

Customizing and Configuring the Solstice Backup Software

Once you have used some of the basic features and functionality of the Solstice Backup software, you can start customizing and configuring Solstice Backup settings for your environment. The Solstice Backup software enables you to configure some or all of its settings. You can also mix customized and preconfigured settings in any combination.

Preconfigured Settings

You can fully use and evaluate Solstice Backup software with the default settings that are automatically assigned to resources and attributes (for example, the Default Pool). Whenever you create a new resource, you always have the option of using default attributes and preconfigured resources.

Configuration Planning

Configuration planning helps you clarify the interdependencies and trade-offs you face in implementing an overall backup strategy. Careful configuration planning is the most critical factor in ensuring that backups run smoothly and that recoveries proceed efficiently and without error.

To plan your configuration:

1. Draw a diagram of the networked computers that will become Solstice Backup servers and clients. Identify potential network trouble spots, such as computers on the other side of a firewall or routers that might not be able to successfully back up to the Solstice Backup server.

2. Estimate the size of the average daily backup window for the enterprise (the period of time set aside for backing up data on the network). Identify special circumstances that might alter the size of the backup window, such as weekend activity, operator availability (if a stand-alone device is being used), or end-of-quarter requirements.

3. Identify the amount of data that must be backed up by each computer and estimate the total amount of backed-up data. Estimate how long backups will take, given the characteristics of the specific environment (for example, network bottleneck speeds, device speeds, and data availability).

4. Organize and prioritize data into clusters that can begin backups at the same time of day (so that they can be organized into Solstice Backup backup groups). This step is most applicable if your backups will require load balancing or staggered backup schedules to accommodate the constraints of the backup window.

5. Decide how long different data should remain in the storage management system for users to browse and recover. The availability of data files might need to be weighed against the space required to maintain the file entries within the client file index.

6. Decide how data should be organized within the storage management system. Using Solstice Backup pools, you can send specific save sets of data to designated collections of volumes. For example, you might want to have one pool for all client backups of your organization's accounting department.

7. Identify the trade-offs and select optimal values for the following performance features:

   ■ Number of devices to use for backups

   ■ Number of devices to reserve for recoveries

   ■ Number of data streams to be interleaved onto a single volume

   ■ Number of data streams to send to each Solstice Backup server and storage node

TABLE 2-3 summarizes the kinds of decisions you need to make prior to configuring Solstice Backup software. The table identifies this information in terms of the information the software needs to perform a scheduled backup, and how that information is expressed in a Solstice Backup configuration.

TABLE 2-3   Configuration Information *(1 of 3)*

Type of Needed Information	How the Information Is Expressed in Solstice Backup Software
When does this backup start?	When the server computer's clock is equal to the time designated as the starting time for a backup group, the Solstice Backup server starts the backup.
	Solstice Backup resource: Group Resource attribute: Start Time
Which clients participate in the backup group?	All clients that are members of an executing backup group are considered candidates for backup. Client data might or might not actually be backed up, depending on the client's schedule. All new clients are, by default, members of the preconfigured group named Default.
	Solstice Backup resource: Client Resource attribute: Group

TABLE 2-3 Configuration Information *(2 of 3)*

Type of Needed Information	How the Information Is Expressed in Solstice Backup Software
How is client data listed for backup?	The Solstice Backup administrator clusters files, directories, and/or filesystems into packages organized by client computer. Each unique package is a *Client resource*; multiple Client resources can be configured for one client computer. A *save set* represents the data backed up for a Client resource during a single backup session.
	Solstice Backup resource: Client Resource attribute: Save Set
Where is client data stored?	Client data is written to the pool of storage volumes that is associated with the client's group. Preconfigured groups have been assigned to preconfigured pools.
	Solstice Backup resource: Pool Resource attribute: Group, Client, Save Set, Level
Which backup level is scheduled for this client's data?	Solstice Backup schedules specify a client's backup level for the day (or whether the backup skips the client that day). Group attributes, if specified, supersede a client schedule. Preconfigured schedules are shipped with Solstice Backup software.
	Solstice Backup resource: Client Resource attribute: Schedule *or* Solstice Backup resource: Group Resource attribute: Schedule *or* Solstice Backup resource: Group Resource attribute: Level
How long will the inventory of backed-up data be available for browsing?	The length of time that client data can be browsed via the nwbackup program is called the *browse policy*. The browse policy is assigned by the administrator for each Client resource. Preconfigured time periods are shipped with Solstice Backup software.
	Solstice Backup resource: Client Resource attribute: Browse Policy
How long will the data be available for recovery?	Data on storage volumes is protected from automatic overwriting for a period that is defined by the *retention policy*. The retention policy is assigned by the administrator for each Client resource. Data is available for recovery indefinitely, until the storage volume is physically relabeled. Preconfigured time periods are shipped with Solstice Backup software.
	Solstice Backup resource: Client Resource attribute: Retention Policy

TABLE 2-3   Configuration Information *(3 of 3)*

Type of Needed Information	How the Information Is Expressed in Solstice Backup Software
Is additional data processing needed?	A Solstice Backup *directive* can be applied to client data to perform commonly requested processing (such as compression and encryption). Preconfigured Directive resources are shipped with Solstice Backup software.
	Solstice Backup resource: Client Resource attribute: Directive
Is additional program execution needed for this environment?	Solstice Backup software can accommodate a customized script supplied by the administrator. For example, your environment might require that the client computer shut down its database before performing a backup.
	Solstice Backup resource: Client Resource attribute: Backup Command
Which users need access to manually back up, browse, or recover data?	Solstice Backup software provides manual backup, browse, and recovery access to all users on the client computer. (To recover data, the user must have read permissions to the data.) Access can be broadened to users on other client computers as well.
	Solstice Backup resource: Client Resource attribute: Remote Access

Solstice Backup Administrator Interfaces

You can administer your Solstice Backup server software by the following methods:

- Graphical Solstice Backup Administrator program (`nwadmin`)
- Character-based Solstice Backup administration interface (`nsradmin`)
- Command line

Solstice Backup Administrator Program

The Solstice Backup Administrator program provides a graphical interface for configuring Solstice Backup servers and clients, assigning clients to servers, creating backup schedules, and managing other information.

It is also possible to perform Solstice Backup administrative tasks by entering Solstice Backup commands at a command line. For more information on Solstice Backup commands, refer to the man page for the command, or refer to the *Legato Command Reference Guide*.

Administrator Privileges

Only users included in the Solstice Backup Administrator list have permission to use and make changes to the Solstice Backup server. Users who do not have root privileges can start and use the Solstice Backup Administrator program, but are not able to make changes to the configuration. To learn more about adding administrators and administrator privileges, see "Server Access Management" on page 96.

TCP/IP Network Requirements

The following are the requirements for TCP/IP network communication using Solstice Backup software:

- The Solstice Backup server name must be added to the domain name system (DNS) file that contains the names and aliases of the servers (hosts) on your network, or to the local hosts file located in /etc.

Note – If you are using DNS, reverse lookup must be correctly configured.

- If the Solstice Backup server is a dynamic host configuration protocol (DHCP) client, it must use a reserved address.
- If you are using DHCP with dynamic addresses, DHCP must be synchronized with DNS.

Solstice Backup Administrator Program Interface

You can perform most Solstice Backup administration tasks in the Solstice Backup Administrator program. The Solstice Backup Administrator program contains the program menus, the Solstice Backup toolbar, and the windows shown in FIGURE 2-2 on page 80.
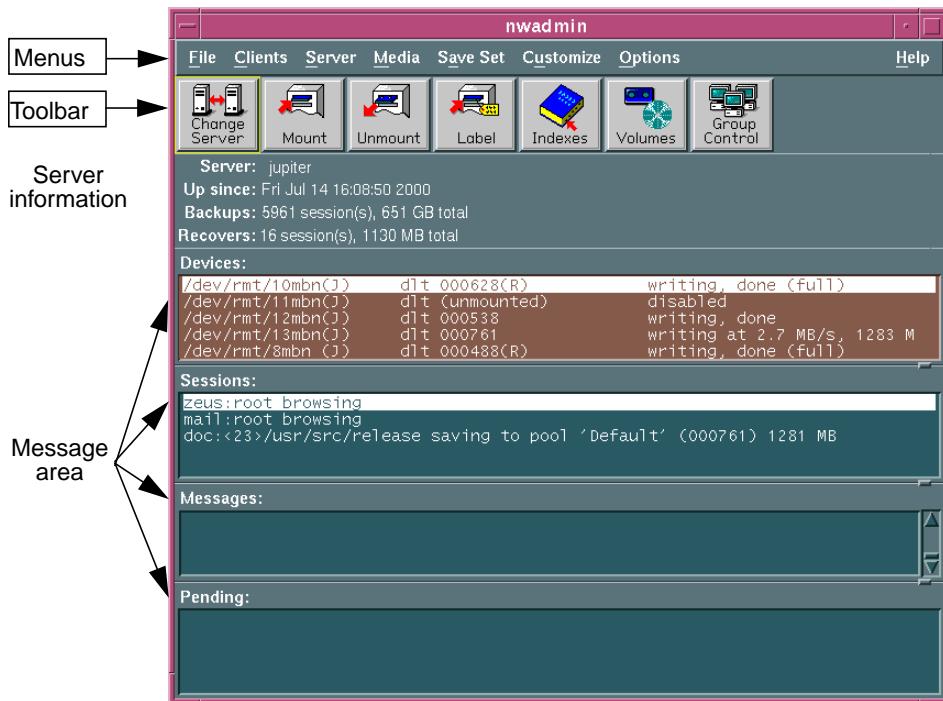
FIGURE 2-2 Administrator Program Interface

Server Information

The server information area displays information regarding the server to which you are currently connected.

Message Area

The message area displays the information listed in TABLE 2-4:

TABLE 2-4 Message Area *(1 of 2)*

Window	Description
Devices window	Displays information about the devices configured for the current server.
Sessions window	Displays progress messages. This information is also written to the daemon.log file in the /nsr/logs directory.

TABLE 2-4   Message Area *(2 of 2)*

Window	Description
Messages window	Displays completion messages. This information is also written to the messages file in the /nsr/logs directory.
Pending window	Displays information on why a backup cannot progress, along with requests to mount tapes and other conditions that require intervention.

The nsrwatch command-line program also provides this information in a character-based format. For more information on the nsrwatch program, refer to the man page for this command, or the *Legato Command Reference Guide*.

Solstice Backup Administrator Toolbar

The Solstice Backup Administrator program's toolbar allows you to perform common tasks by clicking toolbar buttons. The function of each button is described in TABLE 2-5:

TABLE 2-5   Solstice Backup Administrator Toolbar Functions *(1 of 2)*

Button	Name	Function
	Change Server	Opens the Change Server dialog box to connect to a new Solstice Backup server.
	Mount	Mounts the device selected in the Devices window.
	Unmount	Unmounts the device selected in the Devices window.
	Label	Labels the device selected in the Devices window.
	Indexes	Opens the Indexes window, which displays information about the online indexes.

TABLE 2-5   Solstice Backup Administrator Toolbar Functions *(2 of 2)*

Button	Name	Function
Volumes	Volumes	Opens the Volumes window, used to managed volumes.
Group Control	Group Control	Opens the Group Control dialog box, used to monitor groups and manually start and stop group backups.

Customizing the Solstice Backup Administrator Program Interface

You can customize the look and functionality of the Solstice Backup Administrator program.

Hidden Attributes

By default, the Solstice Backup Administrator program leaves some attributes hidden from view in the resource windows. Hidden attributes provide additional functionality and control over Solstice Backup resources.

You can display hidden attributes by selecting Details from the View menu in a particular window.

Foreign Languages in the Interface

Because Solstice Backup software uses a client/server model, and is available in several different languages, there might be times when foreign language messages and other text strings appear in the windows of the Solstice Backup Administrator and other programs. This can occur if you are connected to, or are remotely administering, a Solstice Backup product designed for a different language.

In cases where the foreign language is a double-byte language (for example, Japanese Kanji), your computer must have any appropriate double-byte fonts installed, and must be configured to use the fonts correctly.

Solstice Backup Resources

Each resource has a corresponding dialog box opened by selecting menu commands. Because resource attributes are created and displayed dynamically in the Solstice Backup Administrator program, the same resource dialog box might show different attributes depending on the Solstice Backup server to which you are currently connected.

To find out more about each attribute, select On Window from the Help menu.

For more information on Solstice Backup resources, see "Solstice Backup Resources and Attributes" on page 49.

Character-Based Solstice Backup Interface (nsradmin)

You can use the Solstice Backup character-based interface to perform the same configuration and management tasks as in the Solstice Backup Administrator program.

You can start the character-based interface by entering the following at the command line:

```
# nsradmin
```

For more information about using the character-based interface, refer to the nsradmin man page or the *Legato Command Reference Guide*.

Command-Line Interfaces

You can perform any Solstice Backup configuration or administration task (as well as the tasks available in the client-side backup, recover, archive, and retrieve programs) by entering Solstice Backup commands at the command line. For instructions on how to use these commands, refer to the *Legato Command Reference Guide* or the man pages for individual Solstice Backup commands.

Solstice Backup Manual Backup, Recovery, and Archiving Interfaces

Solstice Backup software uses separate graphical interface programs for performing each of the following tasks:

- Manual backups (`nwbackup`)
- Recoveries (`nwrecover`)
- Archiving (`nwarchive`)
- Retrieval of archived data (`nwretrieve`)

The features and functionality associated with these programs are described in detail in the following chapters:

- Chapter 7
- Chapter 8

The nwbackup Program

Use the `nwbackup` program to perform a manual backup initiated from the client, rather than from the server's scheduled group backup.

To start the `nwbackup` program, enter the following command at the prompt:

```
# nwbackup
```

To specify the Solstice Backup server to which you want to connect, use the `-s` *server* option with `nwbackup`:

```
# nwbackup -s server
```

FIGURE 2-3 on page 85 displays the `nwbackup` program.
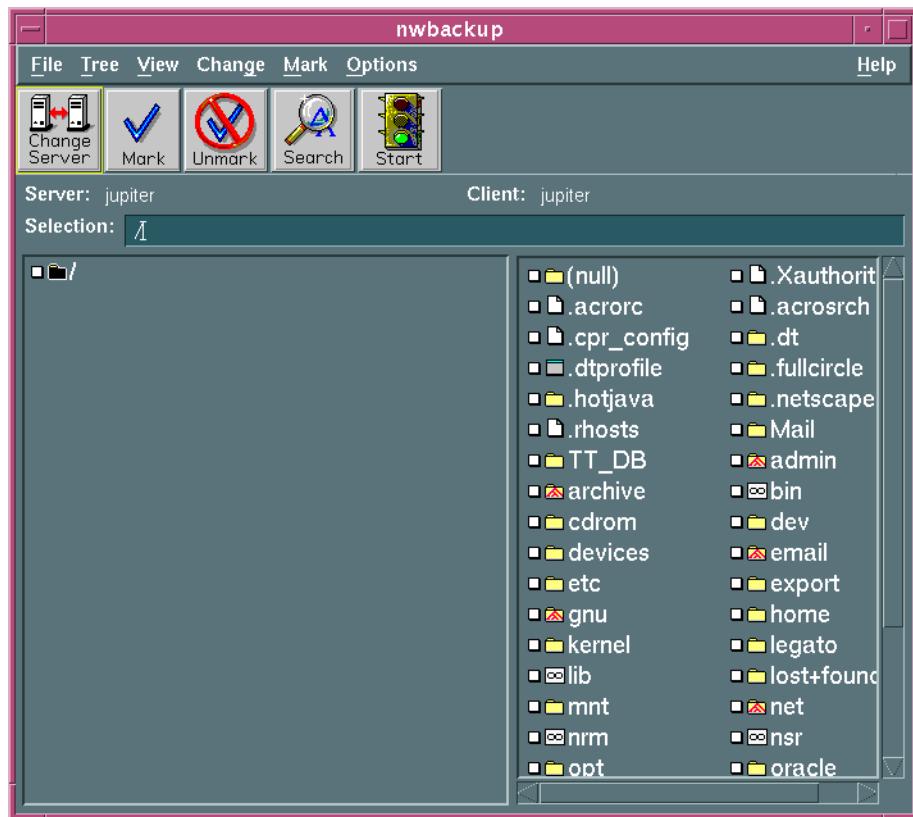
FIGURE 2-3 **nwbackup** Program

The nwrecover Program

Use the nwrecover program to administer and monitor Solstice Backup servers.

To start the nwrecover program, enter the following command at the prompt:

```
# nwrecover
```

To specify the Solstice Backup server to which you want to connect, use the
-s *server* option. To specific the Solstice Backup client, use the -c *client* option:

```
# nwrecover -s server -c client
```

If the -c option is not used, the local computer is used by default.

displays the nwrecover program.

FIGURE 2-4   nwrecover Program

The nwarchive Program

Use the nwarchive program to manually archive files to a Solstice Backup server.

Note – Use of the nwarchive program requires the purchase of a license for the Networker Archive application. You must also enter that licensing information into the Solstice Backup server to fully enable the archiving functionality. See the Solstice Backup Roadmap for licensing information.

To start the nwarchive program, enter the following command at the prompt:

```
# nwarchive
```

To specify the Solstice Backup server to which you want to connect, use the -s *server* option with nwarchive:

```
# nwarchive -s server
```

displays the nwarchive program.
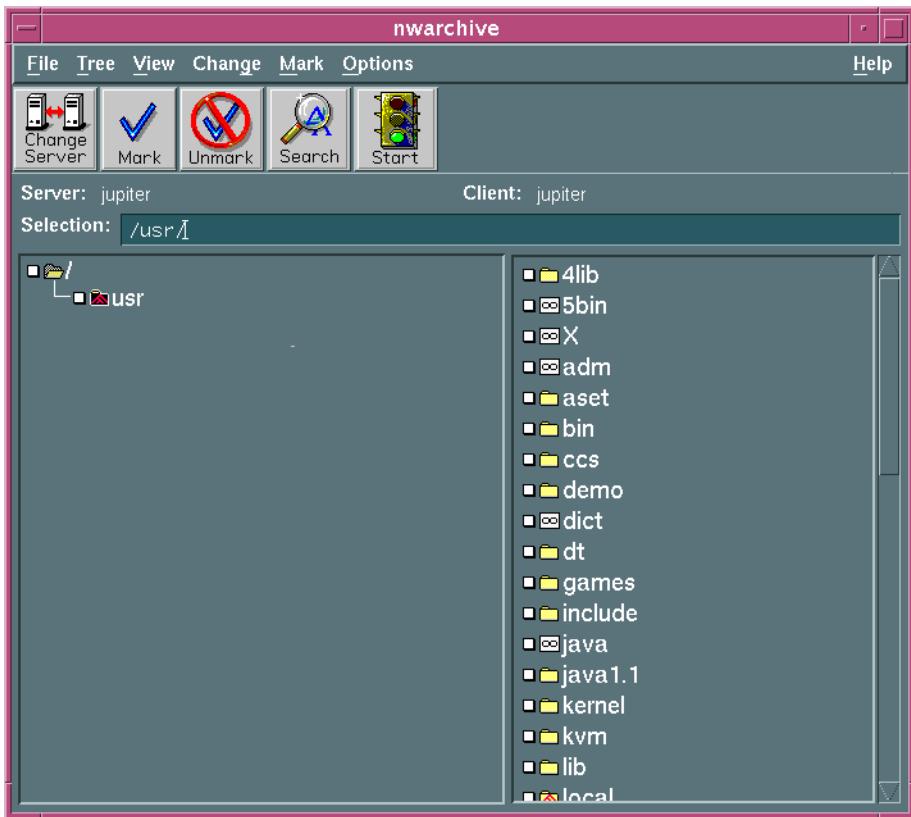
FIGURE 2-5 **nwarchive** Program

For more information on Solstice Backup archiving, see Chapter 8.

The nwretrieve Program

Use the nwretrieve program to retrieve archived files manually from a Solstice Backup server.

To start the nwretrieve program, enter the following command at the prompt:

```
# nwretrieve
```

To specify the Solstice Backup server to which you want to connect, use the
-s *server* option with `nwretrieve`:

```
# nwretrieve -s server
```

displays the `nwretrieve` program.
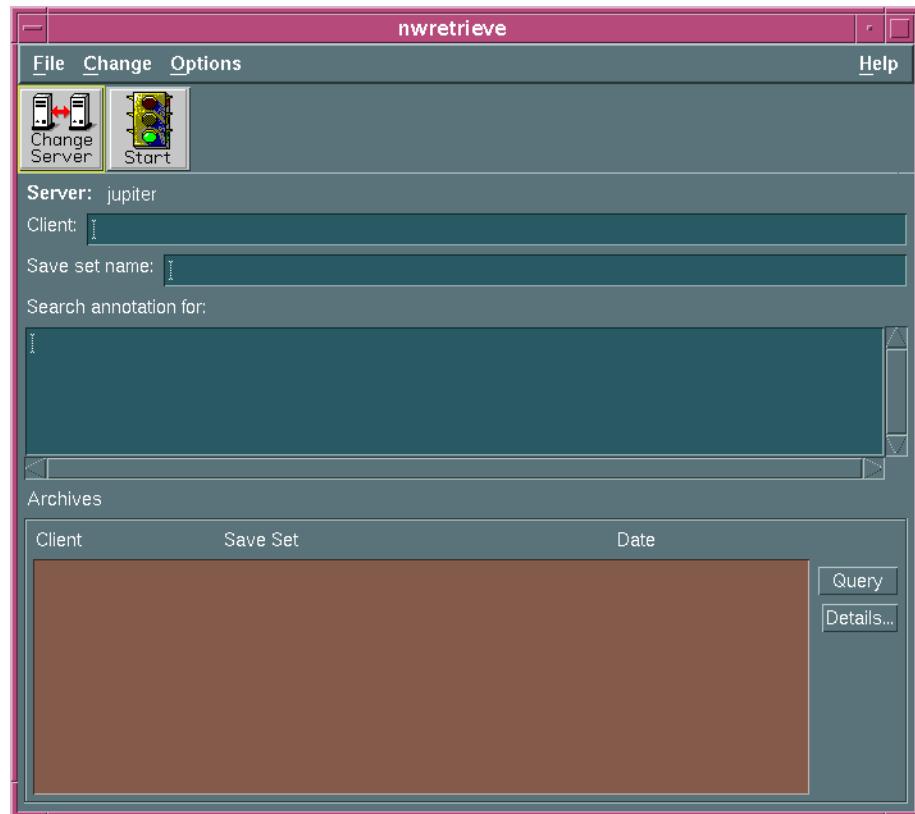
FIGURE 2-6 **nwretrieve** Program

Rehosting a Solstice Backup Server

If you need to transfer the server software to another computer (or transfer the computer to another IP address), you must get a *Host Transfer Affidavit* from Sun Customer Service, complete the form, and submit it to Sun. They will send you new authorization codes, which you must install within 15 days, or the Solstice Backup software becomes disabled.

Caution – Before starting the Solstice Backup software on the new server, you must remove the server software from the computer you transferred it from, or shut down that computer's `nsrd` daemon. If you start the `nsrd` daemon on another computer with the same enabler or authorization code, you will receive a copy violation error message and the software will become disabled.

Server Management

This chapter describes operations that are managed from the Solstice Backup server.

The main sections in this chapter are:

Solstice Backup Server Connectivity

Solstice Backup software allows you to connect to and monitor multiple Solstice Backup servers. The Solstice Backup Administrator program automatically connects to your local Solstice Backup server or to the client's default backup server, if the Solstice Backup Administrator program is launched on a machine with only the Solstice Backup client installed. You can change to another Solstice Backup server at any time in order to configure that server or monitor its activities by following the instructions for "To Establish a Server Connection" on page 92.

Note – Unless you have administrator privileges for that server, you cannot perform any configuration tasks or issue commands from the Solstice Backup Administrator program. You can only monitor the server's activities.

▼ To Establish a Server Connection

To change the Solstice Backup server you are administrating:

1. From the Server menu, select Change to open the Change Server window.

2. From the Server list, select a server. If the server you want is not listed, do one of the following:

   - Click Update Server List to refresh the list.
   - Enter the name of the server in the Server attribute.

3. Click OK.

The Solstice Backup Administrator program now displays information for the new server to which you changed.

▼ To Close the Connection

To quit the Solstice Backup Administrator program and terminate the server connection, select Exit from the File menu.

Client/Server Communication Configuration

Communication between the Solstice Backup server and its clients is defined by configuration attributes entered for the Server and Client resources. How you configure these attributes—such as the level of parallelism or who has administrator privileges—can affect the performance and security of your backups.

For more details about how to set up Solstice Backup clients in the Client resource, see Chapter 4. To diagnose problems with network communications that affect Solstice Backup software performance, refer to the troubleshooting information in "Client/Server Communications" on page 484.

Server Configuration

The following section describes how to configure your Solstice Backup server from the Solstice Backup Administrator program.

▼ To Set Up the Server

To set up your Solstice Backup server:

1. Enter `nwadmin` at the command prompt to start the Solstice Backup Administrator program.

2. From the Server menu, select Setup Server to open the Server dialog box.

3. From the View menu, select Details to display hidden attributes.

4. In the Parallelism attribute, enter the maximum number of save streams that the Solstice Backup server allows to arrive concurrently to the server. For more information on this feature, see "Parallelism" on page 94.

5. For the Manual Saves attribute, select Enabled if you want to allow clients to perform manual backups.

6. For the Public Archives attribute, select Enabled if you want to allow a user to retrieve archived data belonging to another author.

7. For the Update Licenses attribute, select Yes if you want the server to resynchronize its externally managed licenses immediately.

8. In the Administrator attribute, specify all users that should have administrative privileges on this Solstice Backup server, using the format *user@host_name*, for example:

```
root@jupiter
```

For more information, see "Server Access Management" on page 96.

9. For the Volume Priority attribute, select NearLine Priority.

10. Complete the Company Information attributes in order to complete your registration form and obtain your authorization code.

The product serial number is located on the Enabler Certificate you received when you bought your Solstice Backup software. Make sure you enter only the product serial number of the server, not that of any add-on product.

11. Click Apply.

Performance Factor Controls

You can increase backup speed by setting the Solstice Backup server to *multiplex* data on a storage device. The Parallelism and Target Sessions attributes enable data from more than one save set to be written to a single storage volume, using your devices to their maximum capabilities.

Note – By definition, save sets that are multiplexed must belong to the same pool of storage volumes.

Multiplexing optimizes and distributes the flow of data from multiple clients to all of the storage devices available to the Solstice Backup server.

Parallelism

The Solstice Backup *parallelism* feature allows the Solstice Backup server to back up data from many clients in parallel and multiplex the data onto the same backup volume. The parallelism attribute is located in the Server dialog box. For more information on server configuration and the Set Up Server dialog box, see "To Set Up the Server" on page 93.

Parallelism keeps a steady stream of data supplied to the storage device, so that it can operate at higher speeds. Multiplexing also ensures that no particular client can monopolize the server. This means that other clients can back up data to the same server simultaneously.

Each edition of the Solstice Backup software has a different default and maximum parallelism level. Each enabled storage node connected to the Solstice Backup server can increase the parallelism maximum. The maximum parallelism for any Solstice Backup server and storage node combination can vary as described in TABLE 3-1. To learn more about storage nodes, see "Storage Nodes" on page 114.

TABLE 3-1 displays the maximum parallelism levels for this release of the Solstice Backup server product, including the maximum for each enabled storage node.

TABLE 3-1   Maximum Parallelism Levels

Solstice Backup Product	With Each Enabled Storage Node	Without Storage Nodes
Workgroup Edition	Not applicable	Parallelism: 32
	Not applicable	Devices: 4
Network Edition	Parallelism: 32 Maximum = (32 x number of nodes) + 32	Parallelism: 32
	Devices: 16 Maximum = (16 x number of nodes) + 16	Devices: 16
Power Edition	Parallelism: 32 Maximum = (32 x number of nodes) + 64	Parallelism: 64
	Devices: 16 Maximum = (16 x number of nodes) + 32	Devices: 32

By optimizing the efficiency of the server, you can control how many clients can back up their data at one time. You may want to adjust this number of clients under the following circumstances:

- If the server is taking too long to back up all the clients, increase your parallelism setting.
- If the server is so busy backing up clients that it is useless for any other tasks, decrease your parallelism setting.

To disable parallelism entirely, so that only one client can back up to the server at a time, set the Parallelism value to 1.

Target Sessions

The Solstice Backup *target sessions* feature allows you to set the number of backup sessions accepted by an active device. The Target Sessions attribute is configured in the Device resource. Each device connected to a Solstice Backup server or storage node (including each device in an autochanger or silo) can have a different target sessions value. For more information about adding devices, and about the Device dialog box, see "Devices" on page 215.

You can set the Target Sessions value to any number between 1 and 512. This value also determines which device will process the next backup session.

- If a device is already receiving the number of backup sessions determined by the target sessions value, the Solstice Backup server uses the next underutilized device for your backups.

- If all available devices are receiving the number of backup sessions configured by target sessions, the Solstice Backup server overrides the set value and uses the device with the least activity for the backup session.

It is often more efficient for the Solstice Backup server to multiplex multiple save sets to the same device, rather than write each save set to a separate device. For this reason, the Solstice Backup server attempts to assign to each device a number of save sets, up to the value of Target Sessions, before assigning a save set to another device.

Server Access Management

Many Solstice Backup server functions require the operator to have administrator privileges, so all administrators that are authorized to configure the server must have their addresses included in the server administrator list. For more information, see "Administrator Privileges" on page 97.

You must have Solstice Backup Administrator privileges to add, modify, or delete Solstice Backup administrators. Every user or user group listed in the Administrator attribute of the Server resource has permission to modify the Solstice Backup server configuration. If you are attempting to administer a Solstice Backup server and you receive the following error message:

```
user user_name needs to be on administrator's list
```

then you do not have permission to make configuration changes.

▼ To Add an Administrator

Only a current Solstice Backup administrator can add administrators. To add a Solstice Backup administrator:

1. In the Solstice Backup Administrator program, from the server Server menu, select Setup Server.

2. Enter the additional user or group name in the Administrator list box.

3. Click Add.

   The name you added appears in the Administrators attribute.

4. Click Apply.

For details about valid formats for identifying server administrators, see

▼ To Remove an Administrator

Caution – Do not delete *all* administrators, especially root@server_name, because that administrator is required to start Solstice Backup daemons.

To delete a Solstice Backup administrator:

1. In the Solstice Backup Administrator program, from the server Server menu, select Setup Server.

2. Select the name you want to delete in the Administrator list box, then click Delete.

3. Click Apply.

Administrator Privileges

Although any user can view the server's resources from a client computer, only users you specify in the Administrator attribute in the Server resource can add to or change the configuration of the Solstice Backup server, media devices, and autochangers. When you first install the Solstice Backup server, root@server-name is the only user authorized to change the Solstice Backup configuration. The Administrator attribute in the Server dialog box should include the following entries, depending on the platform of the client trying to administer the Solstice Backup server:

- Windows: Administrators@backup-server
- UNIX: root@backup-server
- NetWare: supervisor@backup-server

Note – If a hostname is used, the Solstice Backup server must be able to convert the hostname to a resolved TCP/IP address.

⚠️ **Caution –** If you need to restrict the administrator list so that it does not include `root@server_name`, make sure that the administrator list includes the account that starts the Solstice Backup daemons.

You can give administrative privileges to additional users or user groups by adding their names to the Administrator attribute of the Server resource. User names must be in one of the formats listed in TABLE 3-2 on page 98.

TABLE 3-2 Acceptable Solstice Backup Administrator Formats

Acceptable Formats	Example
local_group@hostname	`Administrators@jupiter`
global_group@hostname	`Domain Admins@jupiter`
root@*hostname*	root@jupiter
account_name@hostname	`anyone@jupiter`
wildcard@hostname	`*@jupiter`
hostname/wildcard	`jupiter/*`
user@domain-name	`me@jupiter.com`
user@hostname	`me@jupiter`
hostname/user	`jupiter/me`
user@wildcard	`me@*`
wildcard/user	`*/me`
user	`me`

⚠️ **Caution –** By using a wildcard before the @ sign, you grant Solstice Backup administrator privileges to any account logged onto that computer, and by using a wildcard after the @ sign, you grant Solstice Backup administrator privileges to any group or account name that precedes the @ sign, regardless of their location.

Using wildcards in the administrator list may introduce security problems; before using them, verify that your enterprise security is not compromised.

Note – If you use the `nsradmin` interface to input these entries, they must be separated by commas.

Notifications

A *notification* provides information about events that occur in your Solstice Backup environment. The information may be sent to an administrator via e-mail, logged to a file, or acted on by other actions, such as sending an SNMP message to an SNMP monitoring agent.

Notifications allow you to configure the events you want reported, as well as how the Solstice Backup server reports them to you. Specific programs can be executed when an event occurs, including third-party programs. By default, the Solstice Backup server sends notifications to log files located in the */nsr/logs* directory, and depending on the notification being sent, the server might also display notifications in a Solstice Backup Administrator program status window.

Preconfigured Notifications

By default, Solstice Backup software is preconfigured to provide most of the notifications required to monitor Solstice Backup events. TABLE 3-3 lists these preconfigured notifications and the associated actions performed by the Solstice Backup server.

TABLE 3-3   Preconfigured Notifications *(1 of 2)*

Notification	Action Performed by the Solstice Backup Server
Bootstrap	Sends the bootstrap information (important disaster recovery information) to the Solstice Backup server's default printer.
Cleaning cartridge expired	Sends e-mail to root: replace the expired cleaning cartridge.
Cleaning cartridge required	Sends e-mail to root: insert a cleaning cartridge.
Device cleaned	Sends e-mail to root: device cleaning is complete.
Device cleaning required	Sends e-mail to root: clean the device.
Device disabled	Sends e-mail to root: a device has been automatically disabled.
Index size	Sends e-mail to root: check the size of the client file index because it will soon exceed the space available.

TABLE 3-3   Preconfigured Notifications *(2 of 2)*

Notification	Action Performed by the Solstice Backup Server
Log default	Directs the UNIX syslog facility (`/usr/bin/logger`) to log information from the `/nsr/logs/messages` file into the appropriate system log file.
Migration attention	Sends e-mail to root: check the status of a migration operation.
Migration completion	Sends e-mail to root: a migration operation is complete.
Registration	Sends e-mail to root: check the registration status.
Savegroup completion	Sends e-mail to root: degree of success in completing scheduled backups, cloning, and archive operations.
SNMP notification request	Sends event notifications to a network management console (this notification only appears if you have purchased and enabled the Legato NetWorker SNMP module).
Tape mount request 1	Requests media be mounted in a device and displays a pending message.
Tape mount request 2	Requests media be mounted in a device and displays a pending message.
Tape mount request 3	Sends e-mail to root requesting that the tape be mounted.

Customizing Notifications

Notifications require three elements:

- Events
- Actions
- Priorities

These elements are described in the following sections.

Events

TABLE 3-4 lists the type of event notifications that trigger an action.

TABLE 3-4   Events

Event	Description
Media	The media needs operator attention to mount or unmount backup volumes.
Savegroup	A backup group has completed backing up.
Index	The index needs attention.
Server	Other server events have occurred (for example, restarting the Solstice Backup server).
Registration	Product registration needs attention.
Device cleaned	The device has been cleaned.
Device cleaning required	The device requires cleaning.
Cleaning cartridge required	Mount the cleaning cartridge.
Cleaning cartridge expired	The cleaning cartridge has expired and needs replacing.

Note – You cannot edit the events for an existing notification. To change, you must delete the notification and re-create it with different settings.

Actions

TABLE 3-5 lists the actions that the server takes after an event notification occurs.

TABLE 3-5   Actions

Action	Description
logger	Uses the UNIX syslog facility (/usr/bin/logger) to log information or send messages.
lp	Prints the notification.
mail	Sends mail to the specified user.
nsrtrap	Sends notifications to an SNMP management console.

Third-party programs can also be used for the action, as long as the programs support reading from *standard input*. For example, you can use a third-party mail program rather than the UNIX mail command.

Priorities

Each Solstice Backup event has a series of associated messages, and each message has an associated priority. The preconfigured notifications have selected priorities based on the importance of the message being sent. For example, the first time the Solstice Backup server sends a mount backup volume request, the priority assigned to the message is Waiting; the priority of the second request is Critical; and the priority of the third request is Alert.

TABLE 3-6 lists the priorities upon which notifications are based.

TABLE 3-6   Priorities

Priority	Description
Info	Supplies information about the server's status.
Notice	Presents important information.
Warning	Presents information about a nonfatal error.
Waiting	Indicates that the server is waiting for an operator to perform a routine task, such as mounting a backup volume.
Critical	Indicates that the server detected an error that should be fixed.
Alert	Indicates a severe condition that requires immediate attention.
Emergency	Indicates that a condition is imminent that might cause the Solstice Backup software to fail, unless corrected immediately.

Note – You cannot edit the priorities for an existing notification. To change, you must delete the notification and re-create it with different settings.

▼ To Create a Notification

To create a notification:

1. In the Solstice Backup Administrator program, select Customize from the Notifications menu to open the Notifications window.

2. From the View menu, select Details.

3. Click Create.

4. In the Name attribute, enter a name for the notification.

5. In the Action attribute, enter the name of the command you want the Solstice Backup server to use. For example, to use the `nsrtrap` program to send a message to a network management console about a volume that needs mounting, enter:

```
/usr/sbin/nsr/nsrtrap -s 1 host
```

6. Complete the other attributes as necessary, then click Apply.

▼ To Edit a Notification

You cannot edit the name or events for an existing notification. To change either of these, you must delete the notification and re-create it with different settings.

To edit a notification:

1. In the Solstice Backup Administrator program, select Customize from the Notifications menu.

2. From the list displayed in the Notifications attribute, select the notification to edit.

3. Edit the Action attribute and click Apply.

▼ To Delete a Notification

Note – You cannot delete any of the preconfigured Solstice Backup notifications.

To delete a notification:

1. In the Solstice Backup Administrator program, select Customize from the Notifications menu.

2. From the list displayed in the Notifications attribute, select the notification to delete.

3. Click Delete.

4. Click OK to confirm the deletion.

5. Click Apply.

Reports

The following sections describe various aspects of the Solstice Backup reporting function.

Routine Data Movement Operations Reports

The degree of success in the completion of scheduled group backups, group cloning, and archive operations is reported to you by the savegrp program through a Savegroup Completion Report. This report is the program triggered by the preconfigured notification Savegroup Completion. The report is sent as an e-mail to root and to the log file in /nsr/logs/messages.

The report consolidates the following information:

■ Success or failure of each save set participating in the operation

■ Operation's save date and time

■ Bootstrap save set ID

■ Bootstrap volume location (volume name, starting record number, and ending record number)

Solstice Backup Server Bootstrap Report

When the Solstice Backup server finishes a scheduled backup that includes the Solstice Backup server, it creates a *bootstrap* save set, which contains backups of the client file index, media database, and Solstice Backup configuration files. The data in the bootstrap save set is required to re-create your Solstice Backup server in the event of a disaster. The bootstrap file is written to media and automatically printed to the default printer. This printed report is the action triggered by the preconfigured bootstrap notification.

Caution – Save your bootstrap report in a safe location. The bootstrap information is essential for recovery from a disaster. It provides a fast way to locate the starting point of a disaster recovery in the event of a server failure. For more information, refer to the *Legato NetWorker Disaster Recovery Guide*.

The nsrinfo program enables you to query the contents of the Solstice Backup client file index. The nsrwatch program enables you to use a character-based interface to monitor Solstice Backup activity as it occurs. For more information, refer to the *Legato Command Reference Guide* and the man pages for these commands.

Storage Management Application Reports

TABLE 3-7 lists the programs that the Solstice Backup software provides to query the contents of the storage management system. For more detailed descriptions of the most commonly used commands and options, refer to the *Legato Command Reference Guide.*

TABLE 3-7   Storage Management Report Programs

Program	Information Reported
mminfo	Contents and mode of the storage volumes and/or the identification numbers and status of the stored save sets.
mmlocate	User-defined location of storage volumes.
nsrinfo	Contents of the client file index.
nsrmm	Status of the storage devices known to the Solstice Backup software.

Solstice Backup Server Statistics and Diagnostic Reports

Messages that report on Solstice Backup diagnostics are displayed in the Solstice Backup administrator interface and are also contained in the Solstice Backup messages file (/nsr/logs/messages). These messages include warning and error conditions and notice of lost connections.

Message Log Files

The messages generated by the Solstice Backup server daemons (nsrd, nsrindexd, nsrmmdbd, and nsrmmd) are contained in the Solstice Backup messages log and the daemon.log file, typically found in the /nsr/logs directory.

Firewall Support

Firewall support enables you to back up Solstice Backup clients and storage nodes that are separated from the Solstice Backup server by a *packet filtering* firewall. A packet filtering firewall looks at each IP packet entering or leaving the network and accepts or rejects it based on user-defined rules.

- A *service port* is used by a Solstice Backup server or storage node to listen for client requests. The default range used for service ports is 7937-9936.
- A *connection port* is used by the Solstice Backup processes to connect to the server during backup and recovery sessions. The default range used for connection ports is 10001-30000.

This section explains how Solstice Backup software uses firewall ports, and provides guidelines on configuring these ports.

Configuring Firewall Support

This section provides guidelines on configuring the service and connection ports used by the Solstice Backup software. This section also provides details on the range of ports required by a server, client, and storage node.

To configure firewall support, do one of the following:

- Use the default settings. For more information, refer to "To Configure Firewall Support Using Default Settings" on page 106.
- Restrict the number of ports the Solstice Backup software uses. For more information, refer to "Restricting the Port Range" on page 107.

If the configured service or connection port range is not large enough, the associated services and processes will not be able to communicate through the firewall. For more information, see "Troubleshooting Port Configuration Problems" on page 114.

▼ To Configure Firewall Support Using Default Settings

To configure the ports on the firewall using the default settings:

1. Set the firewall rules to allow inbound and outbound TCP/UDP packets to the following ports:

    - For service ports, use the range of 7937-9936

- For connection ports, use the range of 10001-30000

⚠ **Caution –** If the release of Solstice Backup Client Software on the client computer or storage node is earlier than Solstice Backup release 5.5, you must set the firewall rules to allow inbound and outbound packets from UDP 111 and 600-1023.

2. Ensure Network Address Translation (NAT) is disabled for each Solstice Backup host. To disable NAT, refer to the documentation accompanying your router.

3. Shut down and restart the Solstice Backup server.

Restricting the Port Range

You need to complete these three procedures when restricting the port range:

1. Determine the minimum port range required by each server, client and storage node. For more information, see "To Determine Minimum Port Range Required" on page 107.

2. Change the default settings on the Solstice Backup server. For more information, see "To Change Solstice Backup Default Settings" on page 111.

3. Complete the configuration of the firewall feature. For more information, see "To Complete the Configuration of Firewall Support" on page 112.

▼ To Determine Minimum Port Range Required

To determine the minimum service and connection port range required by each Solstice Backup server, client, and storage node, see:

- "Server Port Range Requirements" on page 107
- "Client Port Range Requirements" on page 109
- "Storage Node Port Range Requirements" on page 110

To determine how the ports are being used on a specific computer, use the `netstat -a` or `rpcinfo -p` command.

Server Port Range Requirements

This section provides guidelines for determining the minimum number of service and connection ports required by a Solstice Backup server.

- Service Ports

The minimum number of service ports that you can assign to a Solstice Backup server is 10+n, where:

– 10 is the minimum number of services registered on the Solstice Backup server.

   If the number of services registered on the server is greater than 10, nsrexecd will use one port and the additional nine ports can be shared among the remaining services.

– n is the number of instances of the nsrmmd service. An instance of the nsrmmd service is allocated to each tape device configured on the server.

Example: Setting the Server's Service Port Range

If your configuration included four tape devices and 15 registered services, you would need to allocate 14 service ports to the Solstice Backup server. In this case, you could set the service port range for the Solstice Backup server to 8001-8014.

■ Connection Ports

   Typically, a server requires a minimum of 600 connection ports. The minimum number of connection ports that you can assign to a Solstice Backup server is approximately the sum of the following:

– Number of:

   a. Local and remote devices x 10.

   b. Autochangers x 5.

   c. Groups (or instances of savegrp), backing up to the server through the firewall x 5. For each savegrp instance that runs nsrim, an additional three ports will be required.

   d. Solstice Backup Administration tools to be run remotely (such as nwadmin, nwbackup, and nwrecover) + 6.

   e. Noncritical messages being posted. The Solstice Backup server uses UDP ports for posting non-critical events. This number is dependent on the debug level and the number of errors. The larger the debug level set the more messages that will be posted. To turn debugging off, set the debug level to zero.

   f. Number of mechanisms used for copy protection. The Solstice Backup Server uses UDP ports for copy protection. This number is dependent on the copy protection rules specified on the network.

– Parallelism value set in:

   a. Server resource x 15.

b. Client or Server resource used during recovery sessions + 9.

   c. Client or Server resource used during recovery sessions with a storage node + 9.

– Parallelism value assigned to a device x 6, if you are cloning from a storage node to a server.

Example: Setting the Server's Connection Port Range

If the server required 600 connection ports, you could set the connection port range for the Solstice Backup server to 10001-10600.

Using the above server-related examples, you would set the firewall rules to allow inbound and outbound TCP/UDP packets from the following ports:

- Solstice Backup server service ports: 8001-8014
- Solstice Backup server connection ports: 10001-10600
- `nsrexecd`: 7937
- Sun portmapper: 7938

Client Port Range Requirements

This section provides guidelines for determining the minimum number of service and connection ports required by a Solstice Backup client.

- Service Ports

    The Solstice Backup client services do not bind to service ports as specified by the service port range. You do not need to specify a service port range for a client.

    The Solstice Backup client services, `nsrexecd` and Sun portmapper, will bind locally to ports 7937 and 7938, respectively.

- Connection Ports

    Typically, a Solstice Backup client requires a minimum of 100 connection ports. The minimum number of connection ports that you can assign to a client is as follows:

    – Parallelism value specified in the Client resource x 10

    – Number of concurrent operations, such as `nwadmin`, `nwrecover`, and `nwbackup` x 10

Example: Setting the Client's Connection Port Range

If the client required 100 connection ports, you could set the connection port range for the Solstice Backup client to 10001-10100.

Using this example, you would set the firewall rules to allow inbound and outbound TCP/UDP packets from the following ports:

■ Solstice Backup client connection ports: 10001-10100

■ `nsrexecd`: 7937

■ Sun portmapper: 7938

Storage Node Port Range Requirements

This section provides guidelines for determining the minimum number of service and connection ports required by a Solstice Backup storage node.

■ Service Ports

  The minimum number of service ports required by a storage node is equal to the number of tape devices attached to the storage node.

Example: Setting the Storage Node's Service Port Range

If your configuration included four tape devices, then you would need to allocate four service ports to the Solstice Backup storage node. In this case, you could set the service port range for the Solstice Backup storage node to 8001-8004.

■ Connection Ports

  Typically, a storage node requires a minimum of 150 connection ports. The minimum number of connection ports that you can assign to a storage node is approximately the sum of the following:

  – Number of local devices x 10

  – Number of autochangers x 5

  – If cloning from a server to a storage node, multiply the device parallelism x 6

Example: Setting the Storage Node's Connection Port Range

If the storage node required 150 connection ports, you could set the connection port range for the Solstice Backup storage node to 10001-10150.

Using the above examples, you would set the firewall rules to allow inbound and outbound TCP/UDP packets from the following ports:

- Solstice Backup storage node service ports: 8001-8004
- Solstice Backup storage node connection ports: 10001-10150
- `nsrexecd`: 7937
- Sun portmapper: 7938

▼ To Change Solstice Backup Default Settings

You can configure the service and connection ports for the server, client, and storage nodes by changing the assigned default settings using the Solstice Backup Administrator program or `nsrports`.

- To configure the service and connection port ranges using the Solstice Backup Administrator program:

    1. Select Configure Ports from the Options menu.

    2. In the Configure Ports dialog box, enter the hostname or IP address of the computer to which the port configuration applies and click OK.

    3. In the Set Ports dialog box, enter a service ports range and a connection ports range.

    4. Click OK to save your changes.

    5. Click Cancel to exit the Configure Ports dialog box.

- To configure the connection port and service port ranges using `nsrports`, enter the following at the command prompt:

```
# nsrports -s server -a auth_server [-S|-C] range
```

For a description of the options for the `nsrports` command, see

TABLE 3-8 Options for the `nsrports` Command

Option	Description
`-s` *server*	Specifies the system to contact
`-a` *auth_server*	Connects to a remote system that is located on a platform other than the one on which the command is being executed
`-S`	Sets the system's service ports range to the specified range
`-C`	Sets the system's connection ports range to the specified range

▼ To Complete the Configuration of Firewall Support

To complete this configuration process:

1. Set the firewall rules to allow inbound and outbound TCP/UDP packets within the port range that you set on the Solstice Backup server.

2. Set the firewall rules to allow inbound and outbound packets from ports 7937 and 7938. These ports are used respectively by the Solstice Backup services `nsrexecd` and Sun portmapper.

Caution – If the release of Solstice Backup client software on the client computer or storage node, is earlier than Solstice Backup 5.5 you must set the firewall rules to allow inbound and outbound packets from UDP 111 and 600-1023.

3. Ensure Network Address Translation (NAT) is disabled for each Solstice Backup host. To disable NAT, refer to the documentation accompanying your router.

4. Shut down and restart the Solstice Backup server.

Example: Sample rpcinfo Output

The `rpcinfo -p` command can be used to determine port allocations. The sample `rpcinfo` output shown in TABLE 3-9 shows port allocations for five devices in addition to the ten services required by the Solstice Backup server. A single port, 7937, is reserved for the `nsrexecd` program. Three ports, 7938–7940, are shared among nine other services. Five ports, 7941–7945, are allocated for the five `nsrmmd` services. In total, nine ports are allocated for 15 services.

TABLE 3-9   Sample of rpcinfo Output

Program	Version Number	Protocol	Port	Program/Daemon[1]
390113	1	tcp	7937	`nsrexec`
390103	2	tcp	7938	`nsrd`
390109	2	tcp	7938	`nsrstat`
390110	1	tcp	7938	`nsrjb`
390103	2	udp	7939	`nsrd`
390109	2	udp	7939	`nsrstat`
390110	1	udp	7939	`nsrjb`
390107	4	tcp	7940	`nsrmmdbd`
390107	5	tcp	7940	`nsrmmdbd`
390105	5	tcp	7940	`nsrindexd`
390104	105	tcp	7941	`nsrmmd`
390104	205	tcp	7942	`nsrmmd`
390104	305	tcp	7943	`nsrmmd`
390104	405	tcp	7944	`nsrmmd`
390104	505	tcp	7945	`nsrmmd`

1. Programs that register themselves to the portmapper do not have to use their program name as the daemon name. For example, **nsrexecd** registers its daemon as the **nsrexec** daemon rather than the **nsrexecd** daemon.

Troubleshooting Port Configuration Problems

This section provides guidelines for troubleshooting common problems that may occur after the ports have been configured.

- Cannot bind socket to service port in configured range on system *<hostname>*

  The configured service port range is not large enough for the Solstice Backup server services to listen for client requests.

  To resolve this problem:

  1. Increase the service port range configured for the specified host name.

  2. Set the firewall rules to allow inbound and outbound TCP/UDP packets within the port range specified in Step 1.

- Cannot bind socket to connection port in configured range on system *<hostname>*

  The configured connection port range is not large enough for the Solstice Backup backup and recovery sessions to run. The current backup or recovery session has been aborted.

  To resolve this problem:

  1. Increase the connection port range configured for the specified host name.

  2. Set the firewall rules to allow inbound and outbound TCP/UDP packets within the port range specified in Step 1.

Storage Nodes

Storage nodes are host computers with attached *remote storage devices* (devices not physically attached to the controlling Solstice Backup server). A storage node has the physical connection and ownership of the attached devices, but the Solstice Backup server maintains the client file index and media database. Using the Solstice Backup software, you can route client data directly to a storage node's storage devices without the data first going to your Solstice Backup server.

From your Solstice Backup server, you can perform typical storage tasks, such as mounting and labeling volumes for the storage node devices, as well as configuring Solstice Backup resources associated with the storage nodes.

> **Note –** Although any user can view the server's resources from any client computer, only users specified in the Administrator attribute of the Server resource can add to or change the configuration of the Solstice Backup server, media devices, and autochangers.

The following sections contain information about storage nodes:

- "Prerequisites" on page 115
- "Licensing" on page 115
- "Storage Node Configuration" on page 116
- "Timeout Attribute for Storage Node Operations" on page 116
- "Storage Node Affinity" on page 117
- "Storage Nodes and Cloning" on page 118
- "Storage Node Troubleshooting Tips" on page 118

Prerequisites

To operate the Solstice Backup software with storage nodes, the following software must be installed on the storage node. Packages must be installed in the order listed here:

- Solstice Backup client software
- Device driver software for each device
- Solstice Backup storage node software

Licensing

You must obtain and enter the following enabler codes on the Solstice Backup server to use the Solstice Backup software with storage nodes:

- An enabler code for each individual storage node (included on the storage node Enabler Certificate)
- An enabler code for each autochanger or silo connected to each storage node (included in the autochanger and silo Enabler Certificates)

> ⚠️ **Caution –** If you configure and use the Solstice Backup storage node software on a remote computer in evaluation mode, and the evaluation period expires before you install the storage node enabler code, you will not be able to recover backed-up data to the storage node. To recover that data, you will have to move the remote volume to a local device and perform the recovery from there.

Storage Node Configuration

After the required software and enabler codes are installed, you must configure each storage node device. This procedure is similar to the tasks required for other devices, as described in "Devices" on page 215. However, the storage node device name must follow this convention:

```
rd=storage_node_hostname:device_name
```

For example, rd=jupiter.mars.com:/dev/rmt/0hbn

For more details on configuring storage node devices, including those in autochangers and silos, see "Storage Nodes" on page 237.

In order to run jbconfig and scanner on the storage node, the administrator of the storage node needs to be on the server's administrator list. For normal operation, such as after jbconfig is run on the storage node, the administrator on the storage node can be removed from the server's administrator list. After you add the storage node's hostname to the administrator list, one instance of nsrmmd starts on the storage node for each device that it controls.

> ⚠️ **Caution –** You should not attempt to convert an existing Solstice Backup server to a Solstice Backup storage node. There is a potential for conflict when the resource database, media database, and client file indexes from the retired server are merged into databases and client file indexes of the new Solstice Backup server.

Timeout Attribute for Storage Node Operations

By specifying an attribute named "Nsrmmd Control Timeout," which is configured in the Server resource, you can configure the amount of time that a Solstice Backup server waits for a storage node request to be completed. You must have hidden attributes enabled in order to view this attribute.

The Nsrmmd Control Timeout attribute controls how long the server's `nsrd` daemon waits for a storage node request to be completed. If the timeout value is reached without the completion of the request, the operation discontinues and an error message is logged.

The default value assigned to Nsrmmd Control Timeout is 5 minutes; you can specify any value (in whole minutes) within a range of 2 to 30 minutes.

In previous releases of Solstice Backup software, this function was controlled through the NSR_MMDCONTROL environment variable. If `nsrd` initializes on the server and detects that a setting for NSR_MMDCONTROL exists, the following informational message is issued:

```
NSR_MMDCONTROL env variable is being ignored

use 'nsrmmd control timeout' attribute instead
```

If you receive this message, do the following:

1. Shut down the Solstice Backup daemons.

2. Remove the environment setting for NSR_MMDCONTROL.

3. Restart the Solstice Backup daemons.

4. If necessary, start the Solstice Backup Administrator program and adjust the value of Nsrmmd Control Timeout to the setting that was previously assigned to the NSR_MMDCONTROL variable, or to a setting that best meets your current requirements.

Storage Node Affinity

You can determine which Solstice Backup servers and storage nodes receive a client's data—known as *storage node affinity*—by entering their hostnames in the Storage Nodes attribute located in the Client dialog box. The default setting for the Storage Nodes attribute on most Client resources is nsrserverhost (the host Solstice Backup server). If the Client resource of a storage node computer is created after a remote device on the storage node has been created, the default setting of the Storage Nodes attribute is the storage node and the Solstice Backup server.

If a Client resource is created after a storage node is created, and you want the client to back up to that storage node, you must enter the name of the storage node in the client's Storage Nodes attribute *above* the default nsrserverhost. You can add storage node names to this list at any time. The client will direct its data to the first storage node in the list that has an enabled device and is able to receive the data.

Bootstrap Backup on a Storage Node

When the server's index and the bootstrap save set are backed up, the data writes to a device that is local to the Solstice Backup server. A bootstrap cannot be backed up to a remote device, but a *bootstrap clone* can be written to a remote device. If you use `mmrecov` to recover a bootstrap save set, you must recover the data from a local device.

Storage Nodes and Cloning

Storage nodes can clone data in a manner similar to Solstice Backup servers. A hierarchical structure exists between the storage node and the Solstice Backup server to determine where storage node data is routed to be cloned. The following list presents this hierarchy, from highest to lowest:

1. The storage node's Client resource contains the Clone Storage Nodes attribute for cloning data sent to a storage node. Enter the name of each storage node you want to be available for cloning the storage node's data. The storage node will direct its data to the first storage node in the list that has an enabled device and is able to receive the data for cloning.

   Unlike the Storage Nodes attribute, there is no default setting for the Clone Storage Nodes attribute. If there are no storage nodes listed, the Solstice Backup server's settings are used, as described below.

   Note – The Cloning Storage Nodes attribute is only applicable for Storage nodes. It is not used for clients that do not act as storage nodes.

   ---

2. If the Clone Storage Nodes attribute in the storage node's Client resource is not set, the Clone Storage Nodes attribute in the server's Client resource is used.

3. If the Clone Storage Node attribute in the server's Client resource is not set, the Storage Node attribute in the server's Client resource is used.

Storage Node Troubleshooting Tips

If a backup fails, and the following message is displayed:

```
no matching devices; check storage nodes, devices or pools
```

the problem may be related to storage node affinity.

Here are some possible reasons:

- No devices are enabled on the storage nodes in the Storage Nodes list.
- The devices do not have volumes that match the pool required by the backup request.
- All devices are set to read-only.

A common example is when the client has only one storage node in its Storage Node list, and all devices on that storage node are disabled.

You must fix the problem and restart the backup. To fix the problem in this example, do one of the following:

- Enable devices on one of the storage nodes in the client's list.
- Correct the pool restrictions for the devices in the storage node list.
- Add another storage node to the list that has enabled devices meeting the pool restrictions.
- Set one of the devices to read/write.

Online Indexes

The Solstice Backup server tracks the files it backs up in two databases, which are stored on the server's local filesystem:

- The *client file index* tracks the files that belong to a save set. There is one client file index for each client.
- The *media database* tracks the following:
  - Volume name
  - Backup dates of the save sets on the volume
  - Filesystems in each save set

  Unlike the client file indexes, there is only one media database per server.

The Solstice Backup server can automatically control the size of the client file indexes and media database according to the browse policies and retention policies you set. For more details about using browse and retention policies, see "Managing the Data Life Cycle" on page 162.

The following sections provide information about managing online indexes:

- "Characteristics of the Online Indexes" on page 120
- "Automated Index Activities" on page 121
- "Checking Client File Indexes" on page 121

Characteristics of the Online Indexes

The size of an index is proportional to the number of entries it contains. The media database is usually smaller than the client file index, because the media database stores one entry for each volume, while the client file index stores one entry for each file saved on that volume. The Solstice Backup server selects which volume to mount for recovering a file by mapping the saved files to their volumes.

Each entry in the client file index includes the following information for a backed-up file:

■ Filename

■ Number of blocks

■ Access permissions

■ Number of links

■ Owner

■ Group

■ Size

■ Last modified time

■ Backup time

The online indexes grow with each backup, as entries are added for each newly backed-up file and save set. As long as an index entry for a file remains in the client file index, you can recover the file. Over time, the size of these indexes can grow very large.

> **Caution –** If the filesystem that contains the indexes gets full, the Solstice Backup server is unable to access the media database and is thus unable to access and recover data. Unless you configure the server to control the size of the online indexes by using browse and retention policies, they continue to grow until they exceed the capacity of the filesystem.

Solstice Backup software uses browse and retention policies to manage the life cycle of data, and to automatically control the size of the client file index. For more information on policies, see "Browse and Retention Policies" on page 156.

Automated Index Activities

The Solstice Backup server performs the following online index activities:

- Inserts entries in the client file index for each file saved during a backup. For each new backup, the Solstice Backup server acquires more space from the filesystem for the new entries.
- Removes entries and returns disk space to the operating system. The browse and retention policies automatically determine when entries are removed from the index.

You can also remove save sets manually by clicking Remove Oldest Cycle in the Index Save Sets dialog box.

Checking Client File Indexes

Each time the Solstice Backup server starts, it uses nsrck -ML1 to perform a level 1 consistency check on the client file indexes. In some circumstances, this consistency check will not detect corruption in the client file indexes. If you believe an index may be corrupt, you should run a higher level check on the index, for example:

```
# nsrck -L5
```

If the index is still corrupt, recover the index using the procedure outlined in

▼ To View Information about the Indexes

Information about the client file indexes can be used to:

- Determine the size of the indexes
- Reclaim space used by the indexes
- Determine the save sets associated with a client

To view information about the indexes, select Index from the Clients menu to open the Indexes dialog box, which displays index information for all clients of the server. Index information for all clients of the server appears as shown in
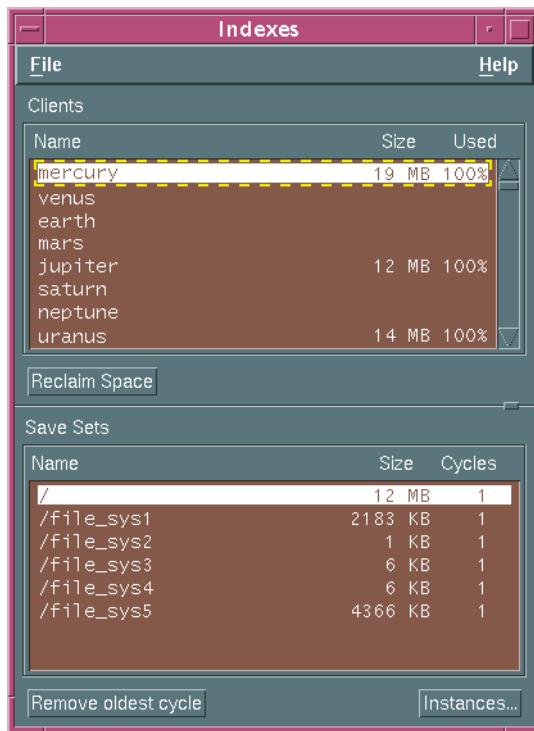
FIGURE 3-1   Indexes Dialog Box

TABLE 3-10 identifies the information displayed for each client in the Clients window of the Indexes dialog box:

TABLE 3-10   Clients Window Description

Column	Description
Client	Names of the Solstice Backup clients that have been backed up by the current server.
Size	Amount of disk space currently allocated to the client file index. As the index size increases, the allocated disk space automatically grows.
Used	Percentage of the index file's allocated space that is currently in use. • If there is no data in the index, 0% is displayed. • If there is data in the index, this value is always 100% because the Solstice Backup software no longer uses disk space reclamation.

Note – The Reclaim Space button is included for backward compatibility with Solstice Backup software prior to version 6.0. For version 6.0 and newer, this feature is no longer necessary.

Index Save Sets

The Save Sets window of the Indexes dialog box displays the save sets assigned to a particular client along with information such as the save set name, size, and number of cycles. The Save Sets window also contains buttons for viewing details on a selected save set and removing old save set cycles.

▼ To View Information about a Client's Save Sets

To view the information about a client's save sets, select the client whose save sets you want to view in the Clients window. The save set information for that client appears in the Save Sets window of the Indexes dialog box.

TABLE 3-11 identifies the information displayed in the Save Sets window for each save set associated with the selected client:

TABLE 3-11   Save Sets Window Description

Column	Description
Name	Name of the save set.
Size	Estimated amount of the index space used by the save set in the client file index.
Cycles	Number of backup cycles available for browsing. A cycle starts with a full backup and ends with the next full backup, and includes any incremental and level 1–9 backups that occur between full backups.

Note – The Remove Oldest Cycle button can be used reduce the size of the client file index by deleting entries for the oldest version of the save set. For more information about reducing the size of your client file indexes, see "Reducing Client File Index Size" on page 127.

Save Set Details

The Instances Details window provides detailed information about a specific save set. Refer to this information to determine how your resources are being used. For example, you can see how large a save set is so you can plan the amount of disk space you need for the online indexes.

▼ To View Save Set Details

To view the save set details:

1. From the Clients menu, select Index to open the Indexes dialog box.

2. Select the client whose save sets you want to view.

3. Click the Instances button to open the Instances window.

FIGURE 3-2 Instances Details Window

TABLE 3-12 identifies the information displayed for each save set instance:

TABLE 3-12   Instances Window Description

Column	Description
Id	Unique identification number of the selected save set.
Files	Number of files backed up during that instance.
Size	Size of the backup.
Date	Date of the backup.
Level	Level of the backup (full, incr [incremental], or level 1–9)

Location of the Client File Index

During the initial client setup, the Solstice Backup software normally designates a default location for the client's client file index on the Solstice Backup backup server. This default location is */nsr/index/client_name*. However, there may be times when you need to designate a different index location when first configuring a Client resource, or you might need to move the file index of an existing client. The following sections address these needs.

▼ To Designate the Client File Index Location

Perform this procedure only when configuring a new Client resource or before the first backup of a new client occurs. If you want to change the location of an existing client's file index, you must move that index. For more information, see "To Move a Client File Index" on page 126.

To designate the client file index location:

1. In the Solstice Backup Administrator program, select Client Setup from the Clients menu to open the Clients window.

2. From the Clients list, select the client.

3. From the View menu, select Details to display hidden attributes.

4. In the Index Path attribute, enter the full path of the directory where the client file index will be stored. If this directory does not already exist, you must create it manually.

5. Click Apply.

6. Navigate to the directory where the client file index for the client resides and enter the following at the command prompt:

```
# nsrck -C client_name
```

If no problems are found when this command is run, then all future client file index information is saved to the designated location.

▼ To Move a Client File Index

In some cases you might want to move a client's file index from its current location to a new location. For example, if the size of the client file index is too large, you can move it to a location with more space.

To move an existing client file index:

1. Run the nsrinfo *client_name* command and redirect the output to a file, for example:

```
# nsrinfo client1 > /tmp/nsrinfo1
```

2. In the Solstice Backup Administrator program, select Client Setup from the Clients menu to open the Clients window.

3. Select the client from the Clients list.

4. From the View menu, select Details to display hidden attributes.

5. In the Index Path attribute, enter the full path of the directory where the client file index will be stored. If this directory does not already exist, you must create it manually.

6. Click Apply.

7. Move the existing client file index to the new location by entering the following at the command prompt:

```
# mv /nsr/index/client_name/* new_index_location
```

8. Run the nsrinfo *client_name* command again and redirect the output to a file. For example:

```
# nsrinfo client1 > /tmp/nsrinfo2
```

9. Verify that the contents of the file create in Step 8 are the same as the contents of the file created in Step 1.

All future client file index information is saved to the new designated location.

Managing the Size of the Online Indexes

Over time, the size of your Solstice Backup server's online indexes can become prohibitively large. You can reduce the size of these indexes using the solutions suggested in the following sections:

- "Reducing Client File Index Size" on page 127
- "Reducing Media Database Size" on page 128
- "To Remove the Oldest Save Set Cycles" on page 128
- "Removing Volume-Based Entries from the Online Indexes" on page 129
- "To Compress the Media Database" on page 131

Reducing Client File Index Size

You can reduce the size of your Solstice Backup server's client file indexes by using one or more of the following procedures:

- Remove save sets that comprise the oldest backup cycle from the client file index. For more information, see "To Remove the Oldest Save Set Cycles" on page 128.
- Remove volume-based entries from the client file index. For more information, see "Removing Volume-Based Entries from the Online Indexes" on page 129.
- Adjust the Browse Policy and Retention Policy attributes of clients backing up to the Solstice Backup server to shorten the period of time that entries remain in the client file indexes. This solution only works for client backups that occur after you change these policy attributes.
- Modify the browse policy associated with a particular save set by using the `nsrmm -w` command. Unless the associated save set contains a large number of files, this method may not be a practical method to reduce the index size. For more information, see "Browse and Retention Save Set Policy Modification" on page 167.

If the size of the client's online file index is still too large, consider moving the location of the index. For more information, see "To Move a Client File Index" on page 126.

Reducing Media Database Size

Reduce the size of Solstice Backup server's media database by using one or more of the following procedures:

- Remove volumes that contain recyclable save sets from the Solstice Backup inventory. For more information, see "Removing Volume-Based Entries from the Online Indexes" on page 129.

  When a volume is removed from the media database, the entries associated with that volume are removed from the media database and the client's online file index. If you select this option, you will still be able to recover the volume's data using the `scanner` program.

Note – You gain very little disk space by removing a media database entry. Leaving a volume's index entries in the media database prevents you from accidentally labeling another volume with the same name.

- Recycle volumes that contain recyclable save sets. For more information, see "To Change the Mode of a Volume" on page 249.

  When a volume is recycled, it undergoes the following procedures:

  - Relabeling
  - Deletion from the media database
  - Reinitialization

  Once a volume is recycled, its contents cannot be recovered.

  To increase the number of currently recyclable save sets, modify the retention policy associated with the current media database by using the `nsrmm -e` command. For more information, see "Browse and Retention Save Set Policy Modification" on page 167.

- Compress the media database. For more information, see "To Compress the Media Database" on page 131.

▼ To Remove the Oldest Save Set Cycles

Client file index entries for a full save set cycle include the last full backup and any dependent incremental or level saves. Removing the oldest cycle frees up disk space. You cannot remove a save set with fewer than two cycles using the Remove Oldest Cycle feature. The last full cycle might be needed for recovery.

To decrease the amount of space used by the index, use the Remove Oldest Cycle feature.

To remove the oldest save set cycles:

1. In the Solstice Backup Administrator program, select Indexes from the Clients menu to open the Indexes dialog box.

2. Select the save sets whose oldest cycle you want to remove, and click Remove Oldest Cycle.

After the Remove Oldest Cycle operation has finished, the statistics in the Indexes dialog box are updated to reflect the current state of the file index.

Removing Volume-Based Entries from the Online Indexes

The main purpose of removing volume-based entries from the online indexes is to eliminate damaged or unusable volumes from the Solstice Backup server. You can also use this feature to reduce the size of the online indexes by purging index entries associated with specific volumes.

Removing Client File Index Entries

You can remove just the entries contained in the client file index using the nsrmm command. This changes the status of the browsable save sets to recoverable:

```
# nsrmm -d -S ssid
```

where *ssid* is the save set ID for the save set. Use mminfo to determine the save set ID:

```
# mminfo -v -c client_name
```

For more information, see the man pages for nsrmm and mminfo or entries for these commands in the *Legato Command Reference Guide*.

When save sets are marked recoverable, users may no longer browse and recover these files using the nwrecover program. You would need to use the save set recover procedure to recover data once the client file index entries are removed. For information on this procedure, see "Save Set Recovery" on page 306.

Removing Client File Index and Media Database Entries

You can remove both the client file index and media database entries. This action removes all knowledge of the volume from the Solstice Backup server. You should only remove a volume from the media database if you think the volume has been physically damaged and is unusable. However, even if you remove the database entries for a volume, as long as the volume is undamaged, the data remains recoverable using the scanner program. For more information on the scanner program, see "The scanner Program" on page 318.

Typically, you do not want to remove both the client file index and media database entries at the same time unless the volume is damaged or destroyed.

Note – The presence of a clone of the particular volume prevents the deletion of the volume entry in the media database. This is because the Solstice Backup server accesses the cloned volume rather than the original volume as needed. The volumes's entry in the media database is never actually purged. Because of this functionality, removing volume entries from the media database is not a particularly effective way to reduce index size.

▼ To Remove Online Index Entries

To remove volume-based entries from the online indexes:

1. In the Solstice Backup Administrator program, select Volumes from the Media menu to open the Volumes window.

2. Select the volume whose entry you want to remove from the online indexes.

3. From the Volume menu, select Remove.

4. Select OK to verify that the volume should be removed.

5. Click Apply.

The Solstice Backup server first cross-checks the indexes before it purges a volume. As a result, the volume might still appear in the Volumes window for a brief period of time.

▼ To Compress the Media Database

You can free up more space on your server by compressing the media database.

To compress the media database:

1. Delete the file */nsr/mm/.cmprssd*.

2. Enter the following command:

```
# nsrim
```

Server Maintenance Tasks

This section describes tasks you might need to perform after you install and configure your Solstice Backup server.

Message Log Management

The Solstice Backup server stores the messages generated by the Solstice Backup server daemons in a message log file in the /nsr/logs directory. When the log file becomes too large, you must delete some messages from the log. You can use environment variables and/or startup scripts to automatically control the size of the log.

▼ To Set the Startup Script to Trim Log Files

To modify the way that Solstice Backup services manage the Solstice Backup log files, set the following environment variables in the Solstice Backup startup script (/etc/init.d/networker) before the lines:

```
(echo 'starting Solstice Backup daemons:') > /dev/console
```

Enter the environment variables using the format:

```
ENV_VAR_NAME = value
export ENV_VAR_NAME
```

The following environment variables can be used to trim log files at startup:

- To change the maximum size of log files, change the NSR_MAXLOGSIZE value. The default value for NSR_MAXLOGSIZE is 1024 KB.

- To change the maximum number of log files that are saved, change the NSR_MAXLOGVERS value. The default value is 4.

For example:

```
NSR_MAXLOGVERS=3
export NSR_MAXLOGVERS
NSR_MAXLOGSIZE=512
export NSR_MAXLOGSIZE
```

Every time the Solstice Backup server starts, it checks the size of the daemon.log file. When the daemon.log file reaches the size defined in NSR_MAXLOGSIZE, or the default size of 1024 KB, the following occurs:

- The daemon.log file is renamed daemon.001 and a new empty daemon.log is created.

- If the daemon.log file fills again, the names of each existing file shift so that the daemon.001 file is renamed daemon.002, daemon.log is renamed daemon.001, and a new empty daemon.log file is created.

- This process is repeated until the value in NSR_MAXLOGVERS is reached, at which point the highest numbered log is removed.

⚠ **Caution –** The trimming mechanism only functions when you start nsrd. The nsrd daemon does not check periodically to see whether the log file has exceeded NSR_MAXLOGSIZE. If nsrd runs for a long time, the log file can still grow very large. To activate the trimming mechanism, enter nsr_shutdown to stop the Solstice Backup daemons, and then restart the nsrexecd and nsrd daemons.

Changing the Name of the Server

Because every Solstice Backup server is also a client to itself, if you change the hostname of the computer on which your Solstice Backup server resides, you must reconfigure the client that resides on the Solstice Backup server in order to back up the server itself.

After renaming the server, all save sets backed up using the old server name will still be identified using that server name. Therefore, you need to retain the Client resource for the old server in order to perform directed recoveries of files that server's backups. For more information, see

After changing the hostname, reconfigure the client that resides on the Solstice Backup server.

▼ To Reconfigure the Server's Client After Changing the Server Name

To reconfigure the server's Client resource after changing the server name and starting the Solstice Backup server:

1. At the command prompt, change to the directory containing the client file index with the old server name, and delete the new server index if it exists in this directory. The default location for the client file index is:

   /nsr/index/old_server_name.domain.com/db6

2. Enter the following at the command prompt to create a new index path for the new Client resource:

   ```
   # nsrck -L2 new_server_name.domain.com
   ```

3. Start the Solstice Backup Administrator program and edit all resource attributes that reference the old server. For example, if the Pool resource has old client name in the Clients attribute, you must change it to use the new client name.

4. Edit the servers file on each client and each storage node associated with this Solstice Backup server to change references from *old_server_name* to *new_server_name*. The servers file is located in the /nsr/res directory.

▼ *To Recover Files Backed Up under the Old Server Name*

To recover files backed up under the old server name, enter the following at the command prompt:

```
# nwrecover -s new_server_name -c old_server_name
```

▼ To Move Your Solstice Backup Server Software to a Different Computer

To move the Solstice Backup server software from one computer to another:

1. Perform a full Solstice Backup backup of all the filesystems on the old Solstice Backup server.

2. Shut down the Solstice Backup daemons on the old server.

3. Move the entire */nsr* directory from the old server to the new server. If */nsr* is a symbolic link on the old server, make sure that the new server has the */nsr* symbolic link set up also.

4. Shut down the new computer, add appropriate hardware devices to the new server, and restart it.

5. Install the Solstice Backup server software on the new server.

Note – If you have an autochanger, do not select the option to start the Solstice Backup daemons. To learn how to install and test the Solstice Backup device drivers, refer to the instructions in the *Installation Guide*.

▼ To Define the Index Entry

Because you created a new host, you must correctly define the index entry for the new host before you start the Solstice Backup daemons. There are two ways to define the index entry:

- Name the new server with the same hostname as the old server at the operating system level before you modify client resources.
- Create a new hostname for the new server with the same configuration choices as the old server.

To create a hostname for the new server:

1. In the Solstice Backup Administrator program, create a new hostname for the new server with the same configuration choices as the old server.

2. Delete the hostname entry for the old server.

3. Shut down the Solstice Backup daemons on both the old and new server.

4. At the command prompt, change to the directory containing the client file index with the old server name, and delete the new server index if it exist in this directory. The default location for the client file index is:

   ```
   /nsr/index/old_server_name.domain.com
   ```

5. Rename the old index directory to the new server hostname, as in:

```
# mv old_hostname new_hostname
```

6. Start the Solstice Backup daemons on the new server.

You see the following messages on the new server:

```
new-server syslog: Solstice Backup Server: (notice) started

new-server syslog: Solstice Backup Registration: (notice) invalid
auth codes detected.

new-server syslog:

new-server syslog: The auth codes for the following licenses enablers
are now invalid.

new-server syslog: The cause might be that you moved the Solstice
Backup server to a new computer.

new-server syslog: You must re-register these enablers within 15 days
to obtain new codes.

new-server syslog:

new-server syslog: License enabler #xxxxxx-xxxxxx-xxxxxx (Solstice
Backup Advanced/10)
```

▼ To Complete the Move

Once you have moved your software, you must reregister your Solstice Backup server. After you move the Solstice Backup server from one system to another, you have 15 days to register the new server with Sun.

To register your new server:

1. Contact the Sun License Center and request a Host Transfer Affidavit.

2. Complete and return the Host Transfer Affidavit to Sun to obtain a new authorization code.

3. Enter the new authorization code into the Auth Code field of the Registration window.

After you have successfully moved your server, complete the following:

1. Verify that all the clients are included in the scheduled backups.

2. Use the nwrecover program to make sure that all the client indexes are visible and therefore recoverable.

3. Back up the indexes on the new server or perform a full backup of the new server as soon as possible.

▼ To Set Up an Old Solstice Backup Server as a Client

If you want to set up the old server as a client, reinstall the Solstice Backup client software according to instructions in your *Installation Guide*.

▼ To Recover Bootstrap Data from Another Server's Volumes Using mmrecov

To recover a Solstice Backup server's bootstrap (media database and resource files) to another server:

1. Move the physical media to the recovering server.

2. Run mmrecov to recover the media database and resource files.

3. Shut down the Solstice Backup server and replace the default resource files with those that were recovered by copying or moving them from /nsr/res.R to /nsr/res.

4. Enter the following at the command prompt to recover the server's indexes:

```
# nsrck -L7 old_server
```

The mmrecov program recovers the media index and server's resource files independent of the server's file index. For more information about the mmrecov command, refer to the *Legato Command Reference Guide* or the mmrecov man page.

Client Operations

Your Solstice Backup product provides preconfigured settings for backing up your clients, including the server, which is considered a client of itself. You can use these preconfigured settings, or create your own. This chapter provides additional information that you might need as you configure your Solstice Backup clients. The information here describes Solstice Backup clients, the associated characteristics of a client and/or a save set, and how to best customize your client configurations to suit the needs of your environment.

The main sections in this chapter are:

Solstice Backup Client

A Solstice Backup client is both a physical computer with Solstice Backup client software installed on it, and a Solstice Backup resource that specifies a set of files and directories that are included in a scheduled backup. As such, a single Solstice Backup client computer can have several individual Client resources specified that back up to the same or different Solstice Backup servers.

The concept of a client computer with multiple Solstice Backup Client resources is further explained in "Multiple Clients from the Same Computer" on page 152.

Solstice Backup client software is available for a variety of platforms, including the most common UNIX and Microsoft Windows operating systems. Solstice Backup products are heterogeneous, meaning that clients and servers of different platforms and operating systems can work together. The basic Solstice Backup client software contains backup and recovery capability for filesystem data; you can purchase additional modules to enable archiving, and backup of a variety of databases.

Client Configuration

This section provides an overview of how Solstice Backup clients are created and integrated into the Solstice Backup storage management environment.

Solstice Backup client software must be installed on computers that are clients to the Solstice Backup server. Because each Solstice Backup server also backs itself up, each Solstice Backup server also has the client software installed. Client resources are created on the Solstice Backup server that specify your configuration choices for each Solstice Backup client.

Client resources determine the following:

■ The data on the client filesystem that will be backed up (the *save set*)

■ The schedule that is followed to back up that data

■ The length of time that the client's backed up data is available for browsing and recovery

■ Any additional features, such as archiving, that are enabled

Client Name

The name that you give the Client resource must be a valid hostname for a computer that has the Solstice Backup client running on it.

If you are creating multiple Client resources for the same computer, you must use the same name for each Client resource. This allows the Solstice Backup server to associate different collections of backed up data with the same client. Unique settings for the Client resource attributes, such as the Save Set attribute, distinguish one Client resource from another.

There are a number of situations where you may want to create multiple Client resources for the same client. For example, if you use a Legato NetWorker Module to back up database files, you may want to create one Client resource to back up the database files and another to back up the client computer's operating system files.

Save Sets

The primary function of a Client resource is to identify the client data to be backed up by the Solstice Backup server. The collection of all data items backed up during a backup session between the server and client is called the *save set*. A save set can consist of a group of files or entire filesystems, as well as application-generated data such as a database or operating system information.

The client's save set is defined in the Save Set attribute of the Client resource. Multiple entries can be included in this attribute. For example, you can enter multiple files or filesystems, or a combination of both:

```
/usr
/home
/etc/hosts
```

If you are backing up large client filesystems, you can use the Save Set attribute to balance load for the client. For example, you could create multiple Client resources, dividing the large filesystems among each resource. You can then create different schedules for each resource to better distribute the client's backup load.

Predefined Save Sets

In addition to entering files and/or filesystems in the Save Set attribute of the Client resource, you can also use the following predefined save sets when configuring your Solstice Backup client:

■ All

- SYSTEM STATE
- SYSTEM DB
- SYSTEM FILES

The All Save Set

The default setting used for the Save Set attribute when a client is first created is All, which refers to the All save set. For UNIX operating systems, the All save set generally backs up all locally mounted filesystems. For Microsoft Windows NT and Windows 2000 operating systems, All also backs up operating system files that determine the client system's state.

TABLE 4-1 provides a list of the files backed up with the All save set.

TABLE 4-1   Components in the All Save Set

Operating System	Files/Save Sets
UNIX	Backs up all filesystems listed in the following locations: • Solaris and Sequent: `/etc/vfstab` • HP-UX, Compaq Tru64, Linux, and SGI: `/etc/fstab` • AIX: `/etc/filesystems`
Windows NT	SYSTEM STATE (includes the registry) REPAIRDISK All fixed disks
Windows 2000	SYSTEM STATE SYSTEM DB SYSTEM FILES All local, physical drives

For more information about the SYSTEM STATE, SYSTEM DB, and SYSTEM FILES save sets, see "Microsoft Windows NT and Windows 2000 SYSTEM Save Sets" on page 140.

Microsoft Windows NT and Windows 2000 SYSTEM Save Sets

Microsoft Windows NT and Windows 2000 contain files and databases used by the operating system to determine the state of your computer at any given time. A simple backup of the computer's filesystem does not capture system state

information. As a result, Solstice Backup software includes three save sets to allow you to back up files that define the system state. TABLE 4-2 provides a list of the components in each of these SYSTEM save sets.

TABLE 4-2   Components in the SYSTEM Save Sets

Save Set	Components
SYSTEM STATE	Includes the Registry, Program Counters and Repairdisk. For Windows 2000, also includes the SYSTEM DB and SYSTEM FILES save sets.
SYSTEM DB	Backs up Windows 2000 system databases
SYSTEM FILES	Backs up Windows 2000 System File Protection components

These save sets can be specified in the Save Set attribute of the Client resource, along with file-based save set specifications. For example, you could enter the following in the Save Set attribute:

```
SYSTEM STATE
C:\
```

This would back up the computer's system state, as well as the "C:" physical drive.

▼ To Create a Client

Before a client can be backed up by a Solstice Backup server, the client computer must have the appropriate Solstice Backup client software installed. Refer to the *Installation Guide* for instructions on installing the client software.

To create a client:

1. In the Solstice Backup Administrator program, select Client Setup from the Clients menu to open the Clients window.

2. Click Create.

3. In the Name attribute, enter the client's hostname. For more information, see "Client Name" on page 139.

4. In the Schedule attribute, select a schedule. For more information, see "Schedules" on page 268.

5. In the Browse Policy attribute, select a browse policy. See "Browse and Retention Policies" on page 156 for more information.

6. In the Retention Policy attribute, select a retention policy. For more information, see "Browse and Retention Policies" on page 156.

7. In the Directives attribute, select a directive. For more information, see "Directives" on page 178.

8. In the Groups attribute, select a group to which your client will be added. For more information, see "Backup Groups" on page 256.

9. In the Save Set attribute, enter the files, filesystems, or predefined save set. Enter All to back up all of the client's filesystems. Multiple entries can be included in this attribute by entering the information and clicking Add. For more information about save sets, see "Save Sets" on page 139.

10. In the Remote Access attribute, enter the user ID or hostnames of other clients that are allowed to back up or recover this client's files. This attribute is used to enable directed recovery. Leave this attribute blank if you do not want to use directed recovery. For more information, see "Directed Recovery" on page 299.

11. In the Remote User attribute, enter the user ID of a client to allow remote access and operations on this client. For Solstice Backup release 6.1 and higher, users who are already defined as administrators of the Solstice Backup server already have remote access to the client and do not need to be included in this attribute.

The Password attribute is only used for clients running NetWare. If you enter a password in this attribute, it is used by the Solstice Backup server to perform remote commands on this client.

12. To instruct the Solstice Backup server to execute customized commands before and after the Solstice Backup server backs up client data, enter savepnc or enter the name of a program file you created in the Backup Command attribute. For more information, see "Backup Command Customization" on page 170.

13. In the Aliases attribute, enter any nicknames for this client computer that queries can match. Otherwise, queries match only the client name.

Note – Each Client resource should include both the DNS short name and long name in the Aliases attribute.

---

14. In the Server Network Interface attribute, enter the hostname of the server's network interface card (NIC) to be used for backups and recoveries. If this attribute is left blank, the NIC associated with the server's hostname is used.

15. In the Archive Users attribute, enter each user who is authorized to use the archive services.

If no users are listed, then only administrators are allowed to use the archive services on the client. The backslash (\) and the at sign (@) are not allowed in the user names for this attribute.

16. In the Storage Nodes attribute, enter the name of each storage node to which the client can back up data.

The client backs up to each active enabled storage node in the order they appear in the attribute. The default storage node *nsrserverhost* represents the server.

17. Enter the storage node names for cloned data in the same manner as Step 16.

18. Complete other attributes as necessary, and click Apply.

Note – If you try to add more clients beyond the number of connections you purchased for your Solstice Backup server, you receive the error message:

```
Too many clients - maximum is number.
```

▼ To Edit a Client

You cannot change the name of a client using this process. To change a client's name, see "To Change a Client's Name" on page 145.

To edit a client:

1. In the Solstice Backup Administrator program, select Client Setup from the Clients menu to open the Clients window.

2. In the Clients list, select the client you want to edit.

3. Edit the attributes of the client, and click Apply.

▼ To Delete a Client

When you delete a client, the Solstice Backup server can no longer back up or recover files from the client computer. The backup history for the client remains in the client file index and media database until the entries are explicitly removed. The browse and retention policies for the remaining entries are replaced with the policies assigned to the Solstice Backup server, defined as a client of itself.

Chapter 4   Client Operations   **143**

Note – If you delete a client from the Solstice Backup server, the client's previously backed-up data is still accessible and recoverable directly from the volume containing the data using the scanner program. For more information, see "Recovering Client File Indexes" on page 316.

To delete a client:

1. In the Solstice Backup Administrator program, select Client Setup from the Clients menu to open the Clients window.

2. In the Clients list, select the client you want to delete.

3. Click Delete.

4. Click OK to confirm that the client should be deleted.

5. Click Apply.

The Client ID

The Client ID attribute of the Client resource allows you to change the name of a Client resource while still maintaining the uniqueness and integrity of that resource. Normally, you would leave this attribute blank when creating or editing a Client resource. Leaving it blank causes the server to automatically assign a value.

The Solstice Backup server determines a value for this attribute based on the following criteria:

■ When creating the first instance of a Client resource, the server creates a new unique identifier.

■ When creating subsequent instances of a Client resource, the new instance inherits the Client ID value from the existing instance of the resource (that is, all Client resources of the same name have the same identification number).

If you create a Client resource and later create a new Client resource using the same client name, the Solstice Backup server recalls the Client ID used previously for this name and reuses it for the new resource.

The only time you would typically need to access the Client ID value of an existing client, or enter a value in the Client ID attribute, is to change the name of a client or to restore a deleted client.

▼ To Access the Client ID

To access the client ID:

1. In the Solstice Backup Administrator program, select Client Setup from the Clients menu to open the Clients window.

2. From the View menu, select Details to display hidden attributes.

3. In the Clients list, select the appropriate client.

The client ID is displayed in the Client ID attribute.

▼ To Change a Client's Name

Note – This procedure is not valid for renaming the client residing on your Solstice Backup server. For instructions about renaming the client that resides on a Solstice Backup server, see "Changing the Name of the Server" on page 132.

To change a client's name:

1. Note the Client ID of the original client. For more information on understanding and accessing the Client ID attribute of the Client resource, see "The Client ID" on page 144.

2. Delete all instances of the original client.

3. Stop all Solstice Backup daemons.

4. On the Solstice Backup server that backs up this client, rename the directory containing the client file index for this client from *old_client_name.domain*.com to *new_client_name.domain*.com. The default location for the client file index is:

   ■ For UNIX:

      /nsr/index/client_name.domain.com

   ■ For Windows:

      %SystemDrive%\Program Files\nsr\index\
      client_name.domain.com

5. Create a new Client resource. For more information, see "To Create a Client" on page 141.

6. In the Client ID attribute, enter the original client's Client ID value.

7. Complete the other attributes as necessary, and click Apply.

8. Restart your Solstice Backup daemons.

Example: Changing a Client's Name

For instance, to change the name from *jupiter* to *mars*, you would remove all instances of client *jupiter* and create the client *mars*, filling in the Client ID attribute of the Client resource with the following value:

74684a4a-00000004-87450361-37e8222d-37e8222d-0001000.

▼ To Restore a Deleted Client

To restore a client that has been deleted from the Solstice Backup server, re-create the Client resource using the same client name and the same attribute values as the deleted client. For information about creating clients, See "To Create a Client" on page 141.

The Solstice Backup server remembers the Client ID for the Client resource bearing this name, and inserts it into the Client ID attribute.

Remote Access Rights for Other Clients

Solstice Backup clients are preconfigured so that only the client itself can browse or recover its own files. However, you can authorize individuals to have remote access rights by entering their user ID in the Remote Access attribute of a client's Client resource. When you enable remote access rights, authorized users can view, back up, and recover files from other Solstice Backup clients.

Note – Users who are included in the Solstice Backup administrator list already have remote access rights to the client and do not need to be included in the Remote Access attribute.

▼ To Grant Remote Access Rights

To grant remote access rights:

1. In the Solstice Backup Administrator program, select Client Setup from the Clients menu to open the Clients window.

2. In the Clients list, select the appropriate client.

3. Enter one of the following in the Remote Access attribute:

   - User ID and hostname (in the format *user@hostname* or *hostname/user*)

   - Group name, if you are using Network Information Service (NIS).

Note – If your company is concerned about security, leave the Remote Access attribute blank, so that only the client itself can recover its backed-up files.

Remote User Attribute

The Remote User attribute is used in the following circumstances:

- NDMP: When using NDMP, enter the NDMP account user name for the computer where the NDMP data module is installed. The user's password must also be entered in the Password attribute.

- Legato NetWorker Modules: When using a Legato NetWorker Module to back up database files, you may need to enter a user name into this attribute to gain access to the application data. For more information, see the `Administrator's Guide` for the appropriate Legato NetWorker Module. The user's password may also need to be entered in the Password attribute.

- NetWare file servers: When using NetWare file servers, the user name to access the files must be entered in this attribute. The user's password must also be entered in the Password attribute.

Solstice Backup Clients of Different Platforms

The Solstice Backup server can back up clients from a variety of platforms. This section provides configuration tips for configuring clients to enable them to back up to the Solstice Backup server.

To use clients of an operating system different from your Solstice Backup server, you must purchase and enable the appropriate ClientPak. See "Solstice Backup Client Licensing" on page 190 for information about ClientPaks and how the Solstice Backup server checks each client before it allows a backup to begin.

Support for 64-bit filesystems exists for clients that run Solaris 7 or later, AIX 4.2, and HP-UX 10.20. You can archive, back up, browse, and recover files larger than 2 gigabytes for clients of Solaris 2.6, AIX 4.2, and HP-UX 10.20. If your clients are not 64-bit capable, you can browse files larger than 2 gigabytes, but you cannot recover them.

UNIX Clients

On all Solstice Backup clients for UNIX, you must manually update and verify certain files and paths, as follows:

■ The */etc/hosts* file must contain the Internet address of the Solstice Backup client and the Solstice Backup server, unless you use DNS or NIS, for example:

```
127.0.0.1 localhost loopback
```

123.456.789.111 server server.domain.com

123.456.789.222 client client.domain.com

The Solstice Backup software does not automatically configure and update the */etc/hosts* file. You must manually edit the file and verify that the information in this file is accurate. Do not delete or "comment out" the entry for the localhost loopback.

- During installation of the SunOS, Solaris, AIX, and DYNIX/ptx client software, if you accepted the default directory when installing the Solstice Backup executables, the default directory should already be in your executable path. If you specified a different directory, add the directory to your executable path for root or Solstice Backup users.

  When you install the HP-UX client software, you must manually add the directory to your executable path, even if you accept the default directory.

  For most UNIX clients, the executable path is set in the PATH environment variable. Adding the directory containing the Solstice Backup executables to your executable path allows execution of Solstice Backup commands without entering the full pathname. For example, you would enter `nwbackup` instead of `/usr/bin/nsr/nwbackup`.

Windows NT Clients

On Solstice Backup clients for Windows NT, you must manually update and verify certain files, directories, and services, as follows:

- If you are not using DNS or Windows Internet Naming Service (WINS), the `%SystemRoot%\Winnt\System32\drivers\etc\hosts` file must contain the Internet address of the Solstice Backup client and the Solstice Backup server.

  Your `%SystemRoot%\Winnt\System32\drivers\etc` directory should contain a sample `hosts` file that gives details about adding entries to the `hosts` file. Do not delete or "comment out" the entry for localhost loopback.

  If you are using DNS or WINS, verify that the DNS or WINS server has entries for both the Solstice Backup client and the Solstice Backup server.

- The `servers` file is typically in `%SystemDrive%\Program Files\nsr\res`. The Solstice Backup software uses the contents of this file to control who has the right to request a program to be executed on this client.

  - If you want this client to back up to other Solstice Backup servers, you must add the names of these servers to this file. You can add only one server name per line.

  - If you want other clients to be able to perform directed recovers to this client, you will need to add their names to the `servers` file. You can add only one client name per line.

  - If you want to allow any Solstice Backup server to back up this Solstice Backup client, delete the `servers` file.

Note – After you save your changes, you must restart the Solstice Backup Remote Exec Service (`nsrexecd.exe`) to make your changes take effect.

To allow any Solstice Backup server to back up this Solstice Backup client, delete the `servers` file.

- The Solstice Backup client for Windows NT must have the latest service pack from Microsoft applied.
- Make sure that the Solstice Backup Remote Exec Service is running.

Note – Solstice Backup Portmapper Service (also known as `portmap.exe`) is an optional service for Solstice Backup clients running software releases prior to 5.7. To enable this service, start it before Solstice Backup Remote Exec Service. The `nsrexecd` process now combines the functionality of `portmapper` and `rpcbind`.

- Save sets maintained on a Windows distributed file system (DFS) link are backed up during a server- or client-initiated backup; however, the Solstice Backup server does not traverse the links or back up the destination files. The server also does not permanently modify the last access time for save sets on a DFS link.

Windows 95/98 Clients

On Windows 95/98 clients, you must manually edit and verify certain files, directories, and services, as follows:

- If you are not using DNS or WINS, the `hosts` file, typically found in `C:\Windows`, must contain the Internet address of the Solstice Backup client and the Solstice Backup server.
- Your Windows 95 directory, typically `C:\Windows`, should contain a sample `hosts` file, named `host.sam`, that gives details about adding entries to an actual `hosts` file. Do not delete or "comment out" the entry for localhost loopback.

  If you are using DNS or WINS, verify that this DNS or WINS server has entries for both the Solstice Backup client and the Solstice Backup server.

- The `servers` file is in `%SystemDrive%\win32app\nsr\res`. The Solstice Backup client software uses the contents of this file to control who has the right to request a program to be executed on this client.
  - If you want this client to back up to other Solstice Backup servers, you must add the names of the additional Solstice Backup servers to this file, one server name per line.
  - If you want other clients to be able to perform directed recovers to this client, you will need to add their names to the `%SystemDrive%\win32app\nsr\res\servers` file, one client name per line.
  - If you want to allow any Solstice Backup server to back up this Solstice Backup client, delete the `servers` file.

- The Windows 95 client must have the latest service pack from Microsoft applied.

- Make sure that the Solstice Backup Scheduled Backup (`wtcpschd.exe`) is running. Put a copy of Solstice Backup Scheduled Backup in the Startup folder to enable scheduled backup to run automatically.

NetWare Clients

On NetWare clients, you must manually update and verify certain files, directories, and services, as follows:

- The `SYS:ETC\HOSTS` file must contain the Internet address of the Solstice Backup client and the Solstice Backup server. The HOSTS file is a simple ASCII text file with one line for each IP address. The IP address is the first entry on the line followed by the hostname and all aliases for each computer. The HOSTS file should also contain an entry for localhost, as in the following example:

```
127.0.0.1 localhost loopback
```

123.456.789.111 server server.domain.com

123.456.789.222 client client.domain.com

Caution – The TCP/IP hostname and the NetWare server name must be identical for the Solstice Backup for NetWare client. In this example, the NetWare server name replaces the value represented by "client."

- TCP/IP must be loaded and bound correctly in the `AUTOEXEC.NCF`, for example:

```
load tcpip
load pcntnw board=1 frame=ethernet_ii name=e_ii
bind ip to e_ii addr=137.69.8.2 mask=255.255.255.000
```

  Load and bind TCP/IP before the Solstice Backup software is installed, or some configuration files are not properly updated.

- Other files that affect Solstice Backup operation on a TCP/IP network and that are automatically configured during the Solstice Backup installation are:

- `SYS:ETC\RPCUSERS`

- `SYS:ETC\SERVICES`

- `SYS:ETC\RPC, SYS:ETC\GATEWAYS`

- `SYS:ETC\NET\NETWARE\SERVICES`

- `RPCNET.CFG` (typically found in `SYS:NSR`)

Other products can also use many of these files, so they might already exist on a client before you install the Solstice Backup software. If the files exist, the Solstice Backup software does not overwrite these files during installation. In most cases, files provided by other software work with Solstice Backup software.

Multiple Clients from the Same Computer

The Solstice Backup server identifies each of its clients by the client computer name. To provide optimal flexibility, the server lets you define multiple Client resources for the same computer with the same computer name, provided that one of the following is true:

- Each client's save sets are unique.

- Clients are included in different backup groups.

- Clients are associated with different schedules.

- Clients are associated with different browse and/or retention policies.

For example, looking at a list of configured Solstice Backup clients, you might see several instances of a client named *mars*, but each instance would contain a unique collection of save sets or would be configured differently regarding groups, schedules, or policies.

Defining multiple clients from the same computer or filesystem can be useful for backing up specialized files, such as databases.

Using Unique Combinations of Clients and Save Sets

Save sets are collections of data to be backed up by the Solstice Backup server. The save sets associated with a specific client instance are visible as a scrollable list in the Save Set attribute of the client's Client resource. Save sets can be any of the following:

- All of the data on a client. This is the default condition indicated by the value All
- Partition on a disk
- Single directory
- Single file
- Raw partition (such as from a database)

If a client has a large volume of data, you may want to schedule the client computer for several, separate client and save set backups. By redefining a large filesystem into multiple client and save set instances, you are able to:

- Automatically back up a large client filesystem.
- Balance the load by avoiding a full backup of the entire filesystem at one time.

You can redefine a single filesystem into a series of separate client and save set instances for the filesystem. You can then associate each client and save set instance with a different backup group and/or a different schedule. Associating different client instances with different backup groups varies the start time of the backups. Staggering the start times in this way may achieve the necessary load balancing needed for your environment.

If different backup start times do not reduce the load adequately, you can associate the different client and save set instances with different backup schedules as well. Recall that a client's schedule determines the level of backup (if any) that is run on a particular day. By using different schedules, you can specify that each client and save set instance runs its full backup on a different day of the week. For more information on schedules, see "How the Solstice Backup Server Uses Schedules" on page 268.

Client and Save Set Combination

The same save set can appear in the Client resource's Save Set attribute for multiple client instances. This characteristic permits you to associate the same save set with more than one group or schedule for backup.

If the default value All appears in the Save Set attribute, all local data for the client computer is backed up according to the group and schedule listed in the Client resource.

To configure the client so that a specific filesystem is backed up, change the value in the Save Set attribute to the filesystem pathname. If you want to define more than one save set on a client, enter each save set (partition, filesystem, or file) on a separate line in the Save Set attribute.

Save Set Status Values

The Solstice Backup server assigns to each backed-up save set a status based on the success of the backup or the age of the save set data. The status of a save set is displayed in the Volumes window.

The save set status changes in the following situations:

- When the save set exceeds its browse policy. For more information about browse policy, see "How the Browse Policy Works" on page 157.
- No save set exceeds its retention policy unless all the save sets that are dependent on the save set exceed their retention policy. For more information about retention policy, see "How the Retention Policy Works" on page 160.
- When an administrator manually changes the save set status.

TABLE 4-3 provides a list of all the possible values for save set status.

TABLE 4-3   Save Set Status Values

Status Value	Meaning	Description
abort	aborted	You aborted the backup for this save set manually or a crash occurred during the operation. This save set is considered immediately eligible for recycling.
brows	browsable	The files in this save set retain entries in the client file index. You can restore all the files using an index-based recover.
inpro	in progress	This save set is currently being backed up.
recov	recoverable	The files in this save set do not have browsable entries in the client file index, and have not passed the retention policy.
recyc	recyclable	The save set and all save sets dependent on this save set for recovery have exceeded their retention policies.
scann	scanned-in	The entry for this save set was restored using the scanner program. The policies were reset to the values that were applied when the save set was created (or, if the save set was created on Solstice Backup servers from releases prior to 6.0, when the save sets were upgraded to the 6.*x* format). For more information, refer to "The scanner Program" on page 318.
susp	suspect	An attempt to recover this save set failed. The recover command could not read all the blocks of the save set. This can occur, for example, when there is a bad spot on the tape.

▼ To View the Status of a Save Set

To view the status of a save set:

1. In the Solstice Backup Administrator, select Volumes from the Media menu to open the Volumes window.

Note – You can also click the Volumes button to open the Volumes window.

---

2. Select the appropriate volume.

The status of save sets on this volume is listed in the Status column of the Save Sets window. For more information, see .

▼ To Change the Status of a Save Set

The Solstice Backup server marks a volume *suspect* if an error occurred while reading the media during backup. You may want to change the status of the save set the server marked as *suspect* to *normal* if you know that the save set data is not really suspect, but that the server had difficulty reading the data for other reasons. For example, perhaps the media drive heads were dirty, causing the server to mark the save set *suspect*.

In some cases, you may want to mark a functioning save set as *suspect* in order to omit it from recoveries. This is necessary if you want to recover data from a cloned volume rather than from the original, because the Solstice Backup server will always look for the original save set during recoveries. Only when the original is *suspect* or nonrecoverable will the server then request a cloned save set and its corresponding clone volume. This process requires that a cloned volume containing the cloned save set be available. If a cloned volume containing the cloned save set is not available, then the *suspect* volume is used.

1. In the Solstice Backup Administrator, select Volumes from the Media menu to open the Volumes window.

2. In the Volumes window, select the appropriate volume.

3. If the status is *suspect*, select Save Set>Change Status>Normal to change the status to *normal*. If the status is *normal*, select Save Set>Change Status>Suspect to change the status to *suspect*.

Browse and Retention Policies

You use browse and retention policies to specify how long data is available for recovery. You can specify browse and retention policies for each client and save set.

How the Solstice Backup Server Uses Browse and Retention Policies

The Solstice Backup server maintains one file index for each client computer (regardless of the number of Client resources configured for it), and one media database that tracks data from all clients and all save sets. Each time a backup is completed, the server creates entries for the backed-up files in the client file indexes. The *media database* stores one entry for each save set and storage volume during each backup operation.

Each client file index is a browsable structure of data from a single client computer. Users can specify anything from a single file to a complete filesystem. The Solstice Backup server can then be directed to reconstruct the data during a recover session to look exactly as it did at a specific time. The information that the client index contains and coordinates enables the server to automatically handle situations. Such situations include assembling data from backups based on levels, and renaming or deleting files or directories. The server uses browse policies to manage the life cycle of data and to automatically control the size of the client file index.

The *browse policy* determines how long files are maintained in the client's file index on the Solstice Backup server. During the period of the browse policy, users can browse backed-up data in the Solstice Backup recover program (`nwrecover`), and select individual files or entire filesystems for recovery. After the browse policy for a file is exceeded, the Solstice Backup server automatically deletes the entry for that file. The server deletes these entries to manage the size of the client index, which can grow rapidly: one entry for each file backed up during each scheduled backup of the client.

The *retention policy* determines how long save sets are maintained in the Solstice Backup server's media database. For at least the period of the retention policy, you can recover a client's backed-up save sets from media:

- No save set is considered recyclable until, at a minimum, it has exceeded its retention policy.
- No storage volume can be relabeled and written over until, at a minimum, all save sets on the storage volume have exceeded their retention policies.

Theoretically, entries for a save set or a storage volume can remain in the media database forever, long after the retention policy has been exceeded. Entries are removed from the media database under these conditions:

■ Storage volume is relabeled

■ Entries are manually deleted

The media database is the structure that tracks the location of save sets on storage volumes. The Solstice Backup server uses the retention policy to manage the longevity of Solstice Backup managed data. Data is recoverable as long as entries exist in the media database; there is nothing to be gained by rushing to delete media database entries. For all these reasons, the media database retention policy does not trigger the automatic removal of media database entries. Instead, the retention policy determines how long an entry for a save set remains protected from being accidentally written over.

How the Browse Policy Works

You can recover a file that has an entry in the client file index through the Solstice Backup recover program (nwrecover). This enables users to browse and mark files, and initiate data recovery. Client file index entries are not necessarily deleted the same day that the browse policy is exceeded. The Solstice Backup server does not remove the entry for a file until all the save sets that are dependent on the file have also exceeded their browse policies. In general, the entries for a full backup that are older than the browse policy are not removed until one backup cycle has passed. This extra time ensures that you can reconstruct a file to any point in time included in the browse policy period.

The following examples demonstrate how a browse policy affects data availability in the client file index. For more information about schedules, see "Schedules" on page 268, and for more information about backup levels, see "Backup Levels" on page 278.

Example: One-Week Browse Policy

In FIGURE 4-1 on page 158, both the backup cycle and the browse policy are set at one week. A backup cycle is the length of time between full backups. Entries for the first full backup on October 1 remain in the client file index until all the dependent incremental and level 5 backups exceed the one-week browse policy. The full backup performed on October 1 is not removed until October 15, when the incrementals and level 5 that depend on the full backup expire.
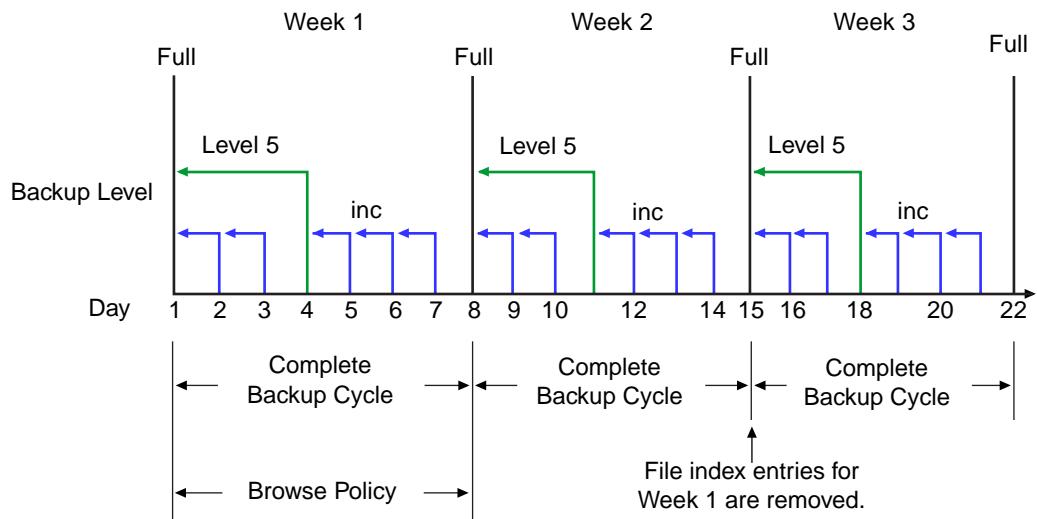
FIGURE 4-1   One-Week Browse Policy

To further illustrate, suppose that on October 12, you decide that you want to
recover information backed up on October 5. The backup performed on the 5th is an
incremental backup dependent on the October 4 backup, which is a level 5 backup.
The October 4 (level 5) backup, in turn, is dependent on the full backup performed
on October 1.

The entry for the full backup performed on October 1 must remain in the client file
index for a period of time equal to the combination of:

- The browse policy (one week)
- One complete backup cycle (one additional week)

That is, until the level 5 backup on October 4 and all incremental backups dependent
on the full backup pass their browse policy. In the example shown in FIGURE 4-1,
entries from the Week 1 backup cycle are removed from the client file index on
October 15.

Example: Two-Week Browse Policy

In FIGURE 4-2, the browse policy is two weeks, which is twice as long as the backup cycle (one week). In this example, on October 18 a user can still find browsable entries in the client file index from backups created on October 4. The backup performed on October 5 is an incremental backup dependent on the October 4 backup, which is a level 5 backup. The October 4 (level 5) backup, in turn, is dependent on the full backup performed on October 1.

The full backup performed on October 1, and the incremental and level backups that depend on it, must remain in the client file index for a period of time equal to the combination of:

- The browse policy (two weeks)
- One complete backup cycle (one additional week)

In this example, entries for the Week 1 backup cycle are not removed from the client index until October 22.
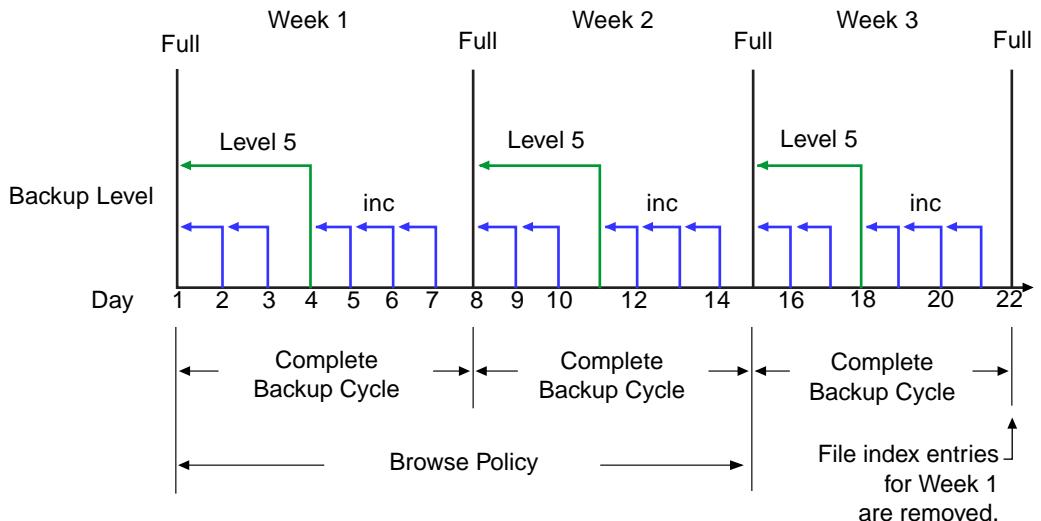
FIGURE 4-2   Two-Week Browse Policy

How the Retention Policy Works

The Solstice Backup media retention policy specifies a period of time during which backed-up data is protected from accidental overwrite. After the retention period is exceeded, the save set is eligible to change its status from *recoverable* to *recyclable*. The term *recyclable* is best understood as "eligible for recycling." The save set's status, however, does not change to recyclable until it, and all the save sets that depend on it, have passed their retention policy. The Solstice Backup server keeps track of save set dependencies regardless of whether the dependent save sets are stored on the same or different volumes. The expiration of a save set's retention policy does not remove the save set's entries from the media database.

The Solstice Backup server changes the mode of that storage volume to recyclable when:

- The retention policy for every save set on a volume expires.
- The status for every save set on a volume changes from recoverable to recyclable.

Since a volume can contain save sets from multiple backup sessions, all with different retention policies, the mode of a volume might not change to recyclable for a long time. All the data on the volume remains available for recovery using either save set `recover` or the `scanner` program. All the entries for recyclable save sets remain in the media database.

The change in status to *recyclable* is a passive reminder that you can overwrite the volume if conditions are right. The volume is available for relabel and use if you:

- Place the volume in an autochanger or mount the volume in a stand-alone device.
- Enable the Auto Media Management attribute in the Device resource.

The existing data is nonrecoverable after the volume is relabeled, so the entries for the overwritten save sets are removed from the media database. For more details about this feature of auto media management, see "How the Solstice Backup Server Selects a Volume for Relabeling" on page 242.

Save set's entries are also removed from the media database when you manually delete a volume from the Solstice Backup volume inventory. However, the data on that volume is still available for recovery using the `scanner` program. The `scanner` program retrieves the information needed to re-create entries in either the client file index, in the media database, or in both places.

- If you re-create the entries in the client file index, a user with the proper permissions can recover data through the `nwrecover` program.
- If you re-create the save set's entries in the media database, a user with Solstice Backup administration privileges can recover data using save set recovery.

For more information on how to use the `scanner` program, see "The scanner Program" on page 318.

Example: Three-Week Retention Policy

FIGURE 4-3 illustrates how a retention policy works. In this example, the backup cycle is set at one week and the retention policy is set at three weeks.

Backup Cycle = 1 week (Full every Sunday)
Retention Policy = 3 weeks
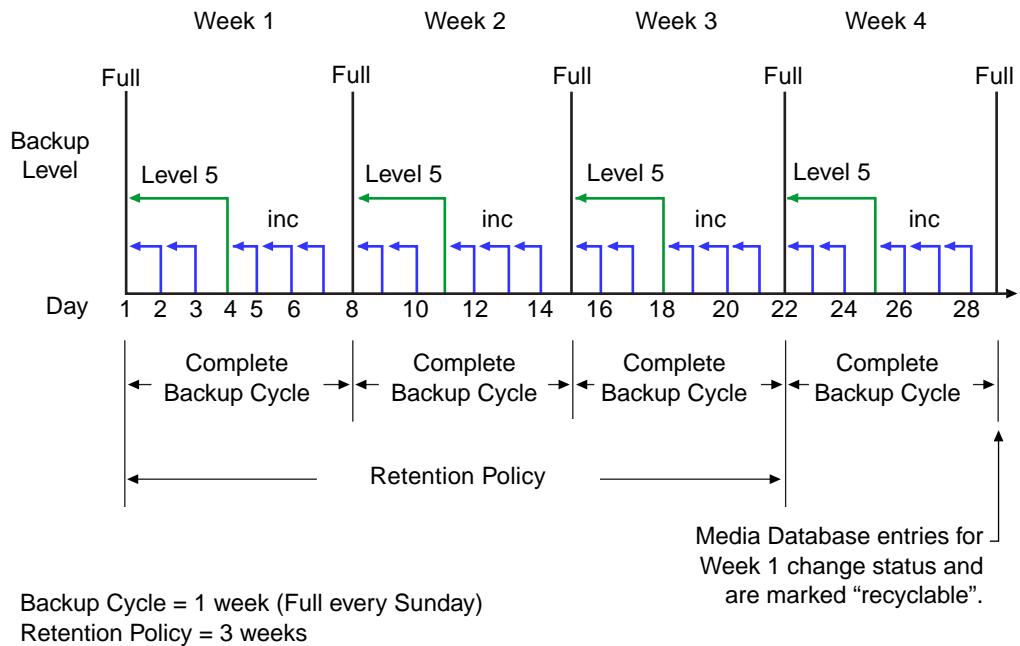
FIGURE 4-3   One-Week Backup Cycle and Three-Week Retention Policy

The save set entries for Week 1 have passed their browse policy and retention policy, but they remain available for recovery using the scanner program until you relabel the volume. When all the save set entries on a volume change status to recyclable, the volume mode changes from full or appendable to recyclable, and the volume is ready to be relabeled for reuse.

Caution – Once you relabel a volume, the data on the volume cannot be recovered.

For more information on these topics, see the following sections:

- Storage volume modes, see "Volume Status" on page 242
- Schedules, see "Schedules" on page 268
- Backup levels, see "Backup Levels" on page 278

Managing the Data Life Cycle

The browse and retention policies that you associate with a client save set control both the growth of the client file index and the media database, and how long data remains available for recovery.

FIGURE 4-4 on page 163 traces the data life cycle through the client file index and the media database. In the example, the entries for the September 1 through September 7 backup cycle remain in the client index for one month (the browse policy), plus the length of a full backup cycle (one week), to ensure that all dependent entries pass their browse policies. In this case, the file index entries for the September 1 through September 7 backup cycle are removed on October 13. Since the entries exist in the client file index, you can browse and recover the data through the nwrecover program. As long as the save set's file entries remain in the client file index, the status of the source save sets is browsable. After the save set status changes from browsable to recoverable, you cannot perform file recovery directly.

The status for each save set backed up during the September 1 through September 7 cycle remains recoverable until their retention policies expire, *plus* however long it takes for all the dependent save sets to pass their retention policies. In this case, the entries from the September 1 through September 7 backup cycle change from recoverable to recyclable on December 8. When all of the save set entries on a volume change status to recyclable, the mode of the volume itself changes from either full or appendable to recyclable.

While the status of a save set is either recoverable or recyclable, you can recover any save set from the storage volume by using either the save set recovery procedure or the scanner program. Alternatively, you can use the scanner program to re-create a save set's entries in the client file index, which enables file recovery directly from the nwrecover program. For more information about using save set recover and the scanner program, see "Save Set Recovery" on page 306 and "The scanner Program" on page 318.

September

S	M	T	W	T	F	S
1f	2	3	4	5	6	7
8f	9	10	11	12	13	14
15f	16	17	18	19	20	21
22f	23					
29f	30					

October

S	M	T	W	T	F	S
		1	2	3	4	5
6f	7	8	9	10	11	12
13f	14	15	16	17	18	19
20f						
27f						

November

S	M	T	W	T	F	S
					1	2
3f	4	5	6	7	8	9
10f	11	12	13	14	15	16
17f						
24f						

December

S	M	T	W	T	F	S
1f	2	3	4	5	6	7
8f	9	10	11	12	13	14
15	16	17	18	19	20	21
22f	23	24	25	26	27	28
29f	30	31				

Backup Cycle = 1 Week

Browse Policy = 1 Month

Retention Policy = 1 Quarter

File index entries from September 1 through September 7 are removed

Save set entries in the media database from September 1 through September 7 change status to recyclable

FIGURE 4-4   Data Life Cycle in the Client Index and the Media Database

On October 13, all data entries from September 1 to September 7 are removed from the client file index. On December 8, the save set entries from September 1 to September 7 in the media database change status from recoverable to recyclable. After all save sets on a volume change status from recoverable to recyclable, the volume mode changes to recyclable. If auto media management is enabled, the volume may be relabeled automatically by the Solstice Backup server to satisfy a volume mount request. After the volume is relabeled, all existing data on the volume is unavailable for recovery. See "Auto Media Management" on page 235 for more information on that feature.

 Caution – When you relabel a volume for reuse within the same pool, the volume identification (the volume name as it appears on the volume label) remains unchanged. Even so, after relabeling, the information that the Solstice Backup server needs to locate and access all existing data on the volume is destroyed, and neither the save set recover feature nor the `scanner` program are options. At this point, the volume is ready for new data. All existing data is inaccessible and is overwritten.

▼ To Assign Multiple Policies to a Single Client

Identical versions of a client and save set combination can have a different set of browse and retention policies assigned for each different backup group to which it belongs. If you create an identical Client resource with the same name and save set values, but assign it to a different backup group, you can designate a different set of browse and retention policies from the original client. The Solstice Backup server employs the Browse Policy and Retention Policy attribute values that correspond to the unique combination of the Client resource's Name, Save Set, and Group attribute values.

Example: Assigning Different Policies for an Identical Client

You already have client *saturn* with a save set value of All and assigned to backup group *General*. The browse policy is Weekly and the retention policy is Monthly. You decide to create an otherwise identical copy of *saturn*, but assign it to backup group *Special*. For this version of *saturn* you designate the browse policy as Weekly and the retention policy as Yearly. If the group *Special* is backed up, then the Weekly and Yearly policies are applied. If the group *General* is backed up, then the policies Weekly and Monthly are used.

Preconfigured Policies

TABLE 4-4 lists preconfigured settings for the Policy resource attributes.

TABLE 4-4　Preconfigured Solstice Backup Policies

Policy Name	Solstice Backup Availability
Decade	Available ten years
Month	Available one month

TABLE 4-4   Preconfigured Solstice Backup Policies *(Continued)*

Policy Name	Solstice Backup Availability
Quarter	Available three months
Week	Available one week
Year	Available one year

▼ To Create a Policy

To create a policy:

1. In the Solstice Backup Administrator, select Customize>Policies to open the Policies window.

2. Click Create.

3. In the Name attribute, enter the name of the policy.

   Choose a name that describes how long the policy lasts. For example, the name *Quarter* indicates the policy lasts a quarter of a year. The name you choose must be unique for the Solstice Backup server.

4. In the Period attribute, select a period.

   This attribute works in conjunction with the Number of Periods attribute. Choose days, weeks, months, or years as the period. For example, a quarterly policy is based on the period of a month (a quarter is three months). Solstice Backup software defines a week as seven days beginning on Sunday, a month as the calendar month, and a year as the calendar year.

5. In the Number of Periods attribute, enter a number.

   This value indicates the number of periods applied to the policy. For example, if you choose *months* for the Period attribute and *three* for the Number of Periods attribute, then the policy lasts for three months (or a quarter).

6. Complete other attributes as necessary, and click Apply.

▼ To Edit a Policy

You cannot change the name of a policy.

The Solstice Backup server does not allow a client's browse policy to exceed its retention policy. This means a file index entry must be removed before the volume on which the file is stored can be marked *recyclable* or removed from the media database.

To edit a policy:

1. In the Solstice Backup Administrator, select Customize>Policies to open the Policies window.

2. From the list displayed in the Policies attribute, select the policy you want to edit.

3. Edit the policy, and click Apply.

▼ To Delete a Policy

You cannot delete any policy currently assigned to a client.

To delete a policy:

1. In the Solstice Backup Administrator, select Customize>Policies to open the Policies window.

2. From the list displayed in the Policies attribute, select the policy you want to delete.

3. Click Delete.

   When prompted, click OK to confirm the deletion.

4. Click Apply.

Browse and Retention Policies for Manual Backups

Browse and retention policies for manual backups adopt the same policies as the instance of a particular client name that has the longest retention time. For example, this can occur if you have three different instances of client *mars*, each with the one of the following retention policy periods:

■ One week

■ One month

■ One year

A manual backup of *mars* with no additional modifications automatically adopts a retention policy of one year and the browse policy that corresponds to the client instance with the one-year retention policy.

You can modify the browse and retention policies of manual backups, but only when performing the backup using the `save` command at the command prompt. See "To Modify a Manual Backup Policy" on page 168 for more information.

For more information on manual backups in general, see "Manual Data Backups" on page 291.

Browse and Retention Save Set Policy Modification

You can modify browse and retention policies prior to a save set backup by editing the Browse Policy and Retention Policy attributes in the Client resource. You can also use the `save` program at the command line to specify policy changes of manual backups.

You can use the `nsrmm` program to modify browse and retention policies inherited by the save set records once the backup has occurred.

The `nsrmm` program updates browse and retention policies using these options:

■ `-e` *retention_time* – updates retention time
■ `-w` *browse_time* – updates browse time

These options (`-e` and `-w`) should be used together with the `nsrmm` option `-S` (to specify save set ID).

Please note that browse time cannot be after retention time, the retention time cannot be before the browse time, and the browse time has to be after the *insertion time*. This is the time this save set record was most recently introduced into the save set database. When the `-e` and `-w` options are used with `nsrmm`, a validity check on these options will fail if any of the following are false:

■ Browse time > insertion time, if `-w` option was specified
■ Retention time > insertion time, if `-e` option was specified
■ Retention time ≥ browse time

Both the browse time (specified with `-w` option) and the retention time (specified with the `-e` option) must be stated in time and date formats acceptable to `nsr_getdate`, which converts them to formats acceptable to the Solstice Backup

server. For details about changing browse and retention time, or about time formats acceptable for use, see the `nsrmm` or `nsr_getdate` man pages, or refer to the *Legato Command Reference Guide*.

Note – With the exception of save time, save sets and volumes are no longer dependent on their corresponding Client resource for their browse and retention policies. This allows more control over whether the save set's file indexes or volumes need to be maintained, deleted, or recycled.

Example: Changing Browse and Retention Policies

These are examples of how one might use `nsrmm` to change browse and retention policies:

1. Change the retention time to midnight, January 1, 2005; change the browse time to midnight, January 1, 2002.

```
# nsrmm -S 3315861249 -e "01/01/05 23:59:59"
-w "01/01/02 23:59:59"
```

2. Change the browse time to six months from the current date and time.

```
# nsrmm -S 5315861249 -w 6 months
```

3. Change the retention time to two years from the current date and time.

```
# nsrmm -S 3315861249 -e 2 years
```

▼ To Modify a Manual Backup Policy

You can modify the browse and retention policies of a manual backup only when performing the backup using the `save` command at the command prompt. You cannot modify manual backups using the graphical `nwbackup` program.

Use the `save` command with the following options in the same manner as modifying scheduled backup policies using `nsrmm`.

- `-y` *retention_time* – updates retention time
- `-w` *browse_time* – updates browse time

For example, if you wanted to perform a manual backup of *file_name* and give it a retention time of one month, you would enter the following command at the prompt:

save -y "1 month" *file_name*

The `nsr_getdate` man page provides the terms and syntax to use when specifying browse and retention times. For more information, see the `nsr_getdate` and `save` man pages or refer to the *Legato Command Reference Guide*.

Reports on Browse and Retention Policies for Save Sets

The `mminfo` command can be used with the `-p` option to display a report on the browse and retention times for save sets. Each line of the report displays the following items:

- Save set creation date
- Stored browse and retention dates (the term *undef* is displayed when connecting to a server using a release of Solstice Backup software earlier than release 6.0)
- Save set ID
- Client name
- Save set name

Granularity

Prior to the 6.0 release of Solstice Backup software, browse and retention policies provided only one form of binding, and that was through the Client resource. This was limiting, given that one policy affected all backups for a Client resource definition for past and future backups. This was also true about modifying policies; it affected all past and future backups.

The Solstice Backup server has the ability to modify browse and retention policies for save sets before and after backups. For example, a user may extend the save sets browse or retention value before the original time limit expires. The recycle functionality of `nsrim` no longer needs to consult the Client resource. Instead, it uses save set browse and retention policies from the save set itself.

The Client resource policy settings provide the initial values for the save set record during backup. Later, these settings can be changed and another backup for the same Client resource can result in save sets having the new policy values. Different save sets for the same Client resource can have different browse and retention policies depending on the values they inherited from the Client resource during backup.

▼ To Override Save Set-Based Browse and Retention Policies

Save sets are retained on volumes and in the media database until the save sets expire. Ordinarily, a save set expires and is recyclable when the save set, and all save sets that depend on it for recovery, exceed their browse and retention policies. However, you can explicitly specify an expiration date for a save set that overrides the browse and retention policies. Dependency rules still apply; this means that a save set is not marked as recyclable until all save sets that depend on it are also marked as recyclable.

Backup Command Customization

You can customize your client backups by creating additional programs (scripts) that affect the way your Solstice Backup server backs up client filesystem data.

The simplest implementation of a customized backup would be to create a special, separate client with a single save set listed in the client's Save set text box (located in the General tab of the Client dialog box). However, you can also create a more elaborate backup program to perform additional tasks.

For example, you can create a program that:

1. Shuts down a mail server or database *before* the Solstice Backup server performs a backup operation.

2. Restarts the mail server or database *after* the backup has completed.

Or you might create a program that:

1. Prints a message (such as, "Backup started at 3:33 a.m.") before the backup operation begins.

2. Executes the backup.

3. Prints a message (such as "Backup completed at 6:30 a.m.") when the backup has completed.

You can customize a client's scheduled backups in either of two ways:

- By creating a script that invokes the save program as part of its instructions, and then entering the name of that script in the Backup Command attribute of the client computer's Client dialog box. When the client is backed up, your customized program is invoked instead of the standard save program. When a client is backed up using save in a customized script, the instructions in your backup program are run separately for each save set backed up for the client. See .
- By entering savepnpc in the Backup Command attribute of the client computer's Client dialog box. This way, the client backup invokes the savepnpc program instead of the save program. The first time the client is backed up, savepnpc creates a default backup program file, which you can then customize for future backups of the client. When a client is backed up using savepnpc, the instructions in the backup program are invoked just once during the backup session. See .

Consider the following questions to determine what level of customization will work best for your environment:

- How much disk space do you have?
- Is there client data that does *not* need to be backed up every time (for example, company e-mail)?
- Do you want to have special messages (in addition to the Savegroup Completion Reports) sent by the Solstice Backup server about the backups it executes?

Using the **save** Command in a Customized Backup Script

You can enter the name of a customized script in the Backup Command attribute in the Client resource that includes additional processing instructions. If an entry exists in the Backup Command attribute, the script associated with it is executed, instead of the default save program, when scheduled backups are initiated.

Caution – Keep in mind that, when you use the save program, a new instance of your customized script is invoked for each save set listed in the client's Save Set attribute, rather than just once for that client (as happens with savepnpc). If you specify a save set value of All, your program is executed for each filesystem on the

client. As a result, if you create a Client resource with a customized backup for a database, a command to shut down the database will be executed for each save set you list.

Note – After you create a customized backup program for a client, immediately try backing up the client. Any configuration or network problems that might prevent a successful backup should show up during this test.

The syntax you use to create the backup program or batch file must adhere to the criteria described in the following list. The list is detailed and includes programming information. Do not attempt to write your own backup command unless you can follow these recommendations.

- The backup program name must begin with either the prefix *save* or *nsr* and must not exceed 64 characters.
- The backup program must reside in the same directory as the Solstice Backup `save` command.
- The Solstice Backup `save` command must be used in the backup program to ensure that the data is properly backed up.
- All commands within the program file must be successfully executed; otherwise, the Solstice Backup server cannot complete the remaining instructions.
- When you invoke the Solstice Backup `save` command, invoke the command with the following arguments: save "$@". Doing so enables the `save` command in your batch file to accept the arguments usually passed to it by the Solstice Backup `savefs` command during a routine backup operation.

▼ To Create a Custom Backup Script Using the `save` Command

To create a custom backup script using the `save` command:

1. Use a text editor to create a script in the directory where the Solstice Backup `save` command resides.

   Commands in this script must be placed in the following order:

   a. Run a preprocessing command before a client backup (optional).

   b. Back up the data using the Solstice Backup `save` command (mandatory).

   c. Run a postprocessing command after a client backup (optional).

2. In the Backup Command attribute of the Client resource, enter the name of the backup program.

3. Back up the client to ensure that the backup command you created works.

Using the **savepnpc** Command in a Customized Backup Script

By using the `savepnpc` command instead of the `save` command, you can run preprocessing and postprocessing commands on a client that execute only once during the client backup, instead of once for each save set. This can be useful if the client is running a database or other program that should be stopped before the client is backed up, and then restarted after the backup has completed. The options for the `savepnpc` command are identical to those for the `save` command. For further details about `savepnpc` command options, see the `savepnpc` command in the *Legato Command Reference Guide*.

To execute `savepnpc`, you must first enter "savepnpc" for the Backup Command attribute in the Client dialog box, and run a backup of the client.

The first time a backup group with a client that uses `savepnpc` runs, a standardized `<group-name>.res` file is created in the `/nsr/res` (UNIX) or `%SystemDrive%\ Program Files\nsr\res` (Microsoft Windows) directory on the client (where `group-name` is the same as the name in the Group resource selected for that client). If the client belongs to multiple backup groups, a separate `<group-name>.res` file is created for each group to which the client belongs. The initial `<group-name>.res` file contains the following default backup type, preprocessing, postprocessing, and timeout attributes:

```
type: savepnpc;
precmd: "echo hello";
pstcmd: "echo bye", "/bin/sleep 5";
timeout: "12:00pm";
```

Once the `<group-name>.res` file exists, you can use your favorite text editor to customize the file's attributes. Your customized instructions are executed the next time the client is backed up.

Before performing a save operation on the client, the modified `savepnpc` command performs any preprocessing commands listed for the precmd attribute in the `<group-name>.res` file. Then it executes the save using the options you specified for the `savepnpc` command itself. After the last save operation is successfully completed, `savepnpc` performs any postprocessing commands listed for the pstcmd attribute.

When editing a `<group-name>.res` file, the following points apply:

- The command environment that is opened by the `savepnpc` command to run a customized backup does not automatically inherit the system's default environment. Specifically, environment variables, including PATH, will either not

exist or will be set to NULL. You must build your environment as part of your preprocessing (precmd) commands, especially the PATH variable. On UNIX clients, be sure to source any needed init files (.profile, .cshrc, .login, etc.).

- If you do not want to include your environment variables in the <group-name>.res file, then you must include full pathnames for all commands and files.

- You must escape any backslash (\) characters in the <group-name>.res file. Therefore, a pathname like c:\mydir\myprogram.exe should be written c:\\mydir\\myprogram.exe. The following is an example of a fully functional <group-name>.res file:

```
type: savepnpc;
precmd: "V:\\usr\\sap\\PDB\\SYS\\exe\\run\\PDB-stop.cmd >
C:\\WINNT\\system32\\PDBStop.log 2>&1";
pstcmd: "V:\\usr\\sap\\PDB\\SYS\\exe\\run\\PDB-start.cmd > C:\\
WINNT\\system32\\PDBStart.log 2>&1";
timeout: "12:00pm";
```

It is not necessary to escape any backslash characters in scripts called in the <group-name>.res. A suggested way to simplify the pathname issue is to put all of your commands in a script or batch file, and then put that script's full pathname on the precmd or postcmd line.

- Text written to standard output by the preprocessing commands will appear in the Solstice Backup completion notices. If this output is verbose, you might prefer to direct the output from the commands in your script to a log file for later examination. Output from the postprocessing commands is thrown away, because the connection to the server has already been dropped when these commands are run. Consider redirecting their output to a log file so you can examine the results if something failed to restart.

- To add more than one command sequence to the precmd and pstcmd attributes, use a comma (,) to separate the commands.

- A complete command line for an attribute must end with a semicolon (;).

- Do not use savepnpc with groups that have a space in the group name.

- On a Windows-based client, do not use "@ECHO OFF" in the <group-name>.res file.

- Resident commands, for which there is no executable file present, like time and dir, will not work as commands in your <group-name>.res file. The log reports that the executable file could not be found.

- If one or more of the following is true:
  - You believe your setup is correct.
  - The backup with savepnpc was working but no longer does.

- A backup for a client is aborted.

- The precmd command fails.

Then look for a file named /nsr/tmp/<group-name>.tmp (UNIX) or %SystemDrive%\Program Files\nsr\tmp\<group-name>.tmp (Windows). If you find this file, delete it. The .tmp file is a lock file that is created when savepnpc runs, and is normally removed when the save sets for a system have been completed.

If a new savegroup session is started and there is an existing <group-name>.tmp file, savepnpc assumes that the preprocessing and postprocessing commands have already run and will skip them; however, it will run the filesystem save sessions.

- The timeout attribute indicates a point in time at which your postprocessing commands are to be run, regardless of whether all of your save sets have been backed up or not. The timeout condition must be specified in a format that nsr_getdate can understand (for details, see the man page for nsr_getdate, **or refer to** the *Legato Command Reference Guide*).

Note – If an invalid time is entered for the Timeout attribute, the timeout is not executed, and there is no error message produced.

Note – The Timeout attribute is optional. If you do not plan to use it, you can comment it out with a pound sign (#) and it will be ignored, as shown here:

```
# timeout: "12:00pm";
```

Note – The value of the Timeout attribute may not be the exact time that postprocessing actually commences. The savepnpc program's pstclntsave subroutine uses a one-minute polling interval to check for the completion of preprocessing tasks. Therefore, the savepnpc.log file (see the sample log below) may show that postprocessing was commenced up to 60 seconds *after* the designated timeout.

▼ To Customize the savepnpc Command for Multiple Groups

To customize pre- and postcommand processing for multiple groups:

1. Copy any existing <group-name>.res file to <new_group-name>.res.

2. Edit the new <new_group-name>.res file.

Message Logging by the `savepnpc` Command

Messages generated by `savepnpc` are written to the `savepnpc.log` file, located in the following location on the client's filesystem:

- UNIX: `/nsr/logs`
- Windows: `%SystemDrive%\Program Files\nsr\logs`

The format of the `savepnpc.log` file is similar to the following:

```
09/03/99 13:56:43 preclntsave: All command(s) ran successfully.

09/03/99 13:57:43 preclntsave: All savesets on the worklist are
done.

09/03/99 13:57:51 pstclntsave: All command(s) ran successfully.
```

Example Script Using the `save` Command

The following is an example script that backs up a ClearCase VOB (version object base). This customized script file must reside in the same directory as the Solstice Backup `save` command (for example, on a Solaris system, the `save` program is installed in the `/usr/sbin/nsr` directory). You enter the name of the script into the Backup Command attribute of the Client resource that is used to back up the ClearCase VOB. As a result, this script is invoked instead of the usual `save` command during a scheduled backup.

Caution – You must include the `save` command within your script in order for the backup to occur. If the script file is not in the same directory as the `save` program, the backup will fail.

The following script locks a ClearCase VOB, performs the backup, and then unlocks the VOB.

```sh
#!/bin/sh

# export the SHELL that we are going to use

SHELL=/bin/sh
export SHELL

# export the correct PATH so that all the required binaries can be
found

case $0 in
/* ) PATH=/usr/atria/bin:/bin:/usr/bin:`/bin/dirname $0`
c=`/bin/basename $0`
;;
* )PATH=/usr/atria/bin:/bin:/usr/bin:/usr/sbin
c=$0
;;
esac
export PATH

# These are the valid statuses which save reports on completion of
the backup

statuses="
failed.
abandoned.
succeeded.
completed savetime=
"

# Perform the PRECMD (Lock VOB)

/usr/atria/bin/cleartool setview -exec
"/usr/atria/bin/cleartoollock -c \
'VOB backups in progress' -vob /cm_data/mis_dev" magic_view >
/tmp/voblock.log 2>&1

# Perform backup on client

save "$@" > /tmp/saveout$$ 2>&1

# cat out the save output

cat /tmp/saveout$$
```

```
# search for the backup status in the output reported by save

for i in ${statuses}; do
            result=`grep "${i}" /tmp/saveout$$`
            if [$? != 0]; then
                  echo ${result}
            fi
done

# Perform the POSTCMD (Unlock VOB)

/usr/atria/bin/cleartool setview -exec
"/usr/atria/bin/cleartoolunlock -vob

/cm_data/mis_dev" \
      magic_view > /tmp/vobunlock.log 2>&

# make sure to gracefully exit out of this shell script

exit 0
```

Directives

Directives are special programs that control how the Solstice Backup server processes files and directories during backup and recovery. A directive contains instructions to assist the backup process, maximize the efficiency of a backup, and handle special files.

Suggested uses of directives include:

- Reducing the amount of data you back up by skipping files and compressing data.
- Instructing the Solstice Backup server to always include a certain file for backup, regardless of the backup schedule that applies to the file.

Note – Directives are not available for Windows 95 or Windows 98 clients.

Local and Global Directives

There are two types of directives: local and global.

- Local directives reside on the Solstice Backup client filesystem and are stored in local directive files named nsr.dir (Microsoft Windows) or .nsr (UNIX). Each directory on the filesystem that requires a local directive must have a local directive file residing in the directory.

- Global directories are resources that reside on the Solstice Backup server and are configured in the Directive resource of the Solstice Backup Administrator program. Directive resources are then specified in the Directive attribute of the Client resource for each Solstice Backup client.

Both local and global directives use the same format, except that global directives contain an additional line indicating the directory on the client computer that is acted upon by the directives. This additional line uses the format <<directory>>.

For example, to skip all *.tmp files in the /export/home/mydir directory, you could create a local directory file, named .nsr, in the /export/home/mydir directory containing the following directive:

```
skip: *.tmp
```

The following global directive would have the same effect:

```
<< /export/home/mydir >>
        skip: *.tmp
```

Application Specific Modules

Directives use Application Specific Modules (ASMs) to process files and directories. ASMs are programs that operate within the Solstice Backup environment to perform various tasks during backup and recovery. For example, the compressasm program instructs the Solstice Backup server to compress data during backup.

ASMs are included in local directive files or server-side Directive resources, and are then executed during back up of client data. Directive files and Directive resources can contain one or more ASM. TABLE 4-5 describes the Solstice Backup ASMs.

TABLE 4-5   Solstice Backup ASMs *(1 of 3)*

ASM Name	Description	Example
`always`	Always backs up the specified files, independent of the change time of the file, and ignores the scheduled level. This ASM can be used to ensure that important client files that change constantly are always fully backed up.	`always: *.c`
`compressasm`	Compresses files so that they use less network bandwidth and fewer volumes. This ASM does not compress directories. The amount of compression achieved is data-dependent. This ASM uses considerable amounts of CPU resources, so its benefits may be limited on low-powered systems. If your storage device compresses data, you can also apply this ASM; however, the additional compression is insignificant.	`+compressasm: .`
`holey`	Handles holes or blocks of zeros when backing up files and preserves these holes during recovery. This ASM is normally applied automatically and does not need to be specified.	`+holey: *`
`logasm`	Instructs the Solstice Backup server to not generate errors when the files specified by this ASM are in use. This ASM is useful for files involved in logging, and other similar files that might change during a backup operation.	`+logasm: *.log`
`mailasm`	Uses mail-style file locking and maintains the access time of a file, preserving "new mail has arrived" flag on most mail handlers.	`+mailasm: *.mbx`
`mtimeasm`	Backs up files using the modification time, rather than the inode change time, to determine which files should be backed up. The modification time is the last time the file's contents were modified, while the inode change time is the last time the file's mode, owner, or link count was changed.	mtimeasm: *.log
`nsrindexasm`	Used to recover from Solstice Backup file index backups performed by Solstice Backup servers prior to release 6.0. During recovery from these older index backups, `nsrindexasm` is invoked automatically by `nsrck` and `mmrecov`.	Not applicable

TABLE 4-5   Solstice Backup ASMs *(2 of 3)*

ASM Name	Description	Example
nsrmmdbasm	Used to process the media database. Normally, nsrmmdbasm is invoked automatically by savegrp and mmrecov, and should not be used in Solstice Backup directives.	Not applicable
null	Does not back up the specified files, but does back up the directory containing the files so entries for the files get added to the online indexes. The Solstice Backup server uses this ASM to back up the online indexes during a scheduled backup.	+null: *.tmp
nullasm	Another name for the null ASM, used for backward compatibility.	See null.
posixcrcasm	Calculates a 32-bit cyclic redundancy check (CRC) for a file during backup. This CRC is stored along with the file and is verified when the file is restored; no verification occurs during the backup itself. With this ASM, it is possible to validate a file at restore time, but it does not provide a way to correct any detected errors.	posixcrcasm: *.?*
pw1	Encrypts and password-protects the specified files. For more information, see "Password ASM" on page 182. **Note:** Available on Microsoft Windows and NetWare versions of Solstice Backup software only.	pw1: *.* password "-o{koo{o"
pw2	Allows for files to be encpted and password-protected using a different password than the one specified in the pw1 ASM. For more information, see "Password ASM" on page 182. **Note:** Available on Microsoft Windows and NetWare versions of Solstice Backup software only.	pw2: *.* password "-o{koo{{"

TABLE 4-5   Solstice Backup ASMs *(3 of 3)*

ASM Name	Description	Example
rawasm	Used to back up /dev entries (that is, block and character-special files) and their associated raw disk partition data. On some systems, /dev entries are actually symbolic links to device specific names. Unlike other ASMs, this ASM follows symlinks, allowing the shorter /dev name to be configured. For more information, including important precautions when using this ASM, see "Important Precautions When Using rawasm" on page 183.	rawasm: /dev/oracle1
skip	Omits the specified files and directories from the backup, and does not place the filename in the online index.	+skip: *.tmp
swapasm	The swapasm does not back up actual file data, but re-creates a zero-filled file of the correct size on recovery. This ASM is used on systems where the swapping device is a swap file that must be recovered with the correct size, but the contents of the swap file are not important and do not need to be backed up or restored.	swapasm: compression.doc

Password ASM

The password ASMs (pw1 and pw2), available for Microsoft Windows and NetWare clients only, are used to encrypt and password-protect data. These directives require a local password file, Networkr.cfg, located in the root directory of the filesystem on which the Solstice Backup software has been installed. The password file contains an encrypted password that is created in the Solstice Backup User program. For more information about entering a password in the Solstice Backup User program, see the *Legato NetWorker Administrator's Guide, Windows Version*. If the password file does not exist, you are prompted to enter a password when restoring the file.

The format for this ASM requires that the password be specified, in encrypted mode, on the line directly below the ASM specification:

```
pw1: *.*
        password "-o{koo{o"
```

To include the encrypted password in the directive, copy the encrypted form of the password from the password file.

Important Precautions When Using rawasm

The rawasm directive is used to back up raw disk partitions. Raw partitions are often used to store active database management system data. If your raw partition contains data managed and updated by an active DBMS product, rawasm alone will not give a consistent backup. The database must not be updating the data in an uncontrolled fashion while rawasm saves or recovers data on the partition. The partition must be offline, the database manager shut down, or the partition placed in an appropriate state for backup. As a result, it is better to use a Legato NetWorker Module application to back up databases.

Similarly, if rawasm is used to save a partition containing a UNIX filesystem, the filesystem must be unmounted or mounted read-only to obtain a consistent backup.

Using rawasm with Recovery

When recovering, rawasm requires that the filesystem node for the raw device exist prior to the recovery. This protects against the recovery of a /dev entry and the overwriting of data on a reconfigured disk. You can create the /dev entry, having it refer to a different raw partition, and force an overwrite if needed. If you create the /dev entry as a symbolic link, the data is recovered to the target of the symbolic link.

Recovery of a raw partition must occur on a system configured with the same disk environment and same size partitions as the system that performed the backup. If the new partition is smaller than the original partition, the recovery is not completed successfully. If the new partition is larger than the original partition, only the amount of data originally saved is recovered, and since rawasm does not discover the size completed, the estimated size reported on recovery is not accurate.

Directive Format

Both global and local directives use a file format specified in the nsr(5) man page. This format contains three primary types of directives:

- Directory specifications
- ASM specifications
- Save environment keywords

Comments can be included in directives as well. Any text after a pound sign (#) is treated as a comment.

Directory Specifications

Directory specifications are used by global directives only. The format for a directory specification is:

```
<<directory>>
```

Where *directory* represents the highest-level directory to which the directives following this entry apply. For example, to apply ASMs to the */etc* directory, you would use the following format:

```
<</etc>>
        directive
```

Multiple directory specifications can be included in the same Directive resource:

```
<</etc>>
        directive
<<</tmp>>
        directive
```

When multiple directory specifications are used, directives that follow a directory specification are applied to that directory until the next directory specification. You cannot use wildcards in directory specifications.

Note – For directives applied to clients on Windows systems, all pathnames must be written in uppercase, and if there is a colon in the path name, the entire path must be enclosed in quotation marks:

```
<<"C:\MYDIR">>
```

ASM Specifications

ASM specification directives specify an ASM to run, as well as parameters associated with the ASM. The syntax for ASM specifications is:

```
[+]asm:  argument
```

Where the optional plus sign (+) indicates that the directive applies to both the current directory and all subdirectories, *asm* is the name of the ASM to be executed, and *argument* is a list of filenames that are acted upon by the ASM. The *argument* can include multiple filenames, separated by spaces, and can use wildcards as well. Only filenames and wildcards are allowed; no directory information can be included. For a description of available ASMs, see "Application Specific Modules" on page 179.

Note – For directives applied to clients on Microsoft Windows systems, all filenames must be written in lowercase.

Using Wildcards in ASM Specifications

Wildcards can be used in ASM specifications to replace a single character or string of characters. Standard filematching patterns are supported.

On UNIX systems, when applying a directive to all files, including hidden files (.* files), use *   .?*. Note that there must be a space between the first asterisk and the dot.

Save Environment Keywords

Save environment keywords are used to control how the current ASM and subsequent ASMs that apply to the current directory and subdirectories will be applied. There are three save environment keywords:

TABLE 4-6   Save Environment Keywords *(1 of 2)*

Keyword	Instruct	Example

forget	Instructs the Solstice Backup server to no longer apply inherited directives (those directives that begin with a +).	The following example skips all *.o files in the directory /usr/src, expect those in the /usr/src/sys directory:
		
		`<</usr/src>`
		`+skip: *.o`
		`<</usr/src/sys>>`
		`forget`
		
		This example uses the skip ASM to instruct the Solstice Backup server to skip all files named *.o in the /usr/src directory and all subdirectories. It then uses the forget keyword to instruct the server to not apply the skip ASM to the sys subdirectory.

TABLE 4-6 Save Environment Keywords *(2 of 2)*

Keyword	Instruct	Example

`ignore`	Instructs the Solstice Backup server to ignore all directives applied to the subdirectories below the current directory.	The following example overrides any local directives set in users' home directories, you would use the following:
		
		`<</home>>`
		`ignore`

`allow`	Used in subdirectories that currently have the `ignore` keyword applied to them, and overrides the ignore.	Building on the above example for the `ignore` keyword, you can allow directives in the `/home/doc` directory to be applied:
		
		`<</home>>`
		`ignore`
		`<</home/doc>>`
		`allow`

Preconfigured Directive Resources

The Solstice Backup software comes with a number of preconfigured Directive resources. Some of these Directive resources contain preset global directives, others do not contain preset directives but can be modified to fit the needs of your environment.

All preconfigured Directive resources can be modified, but they cannot be deleted.

TABLE 4-7 lists the preconfigured directives and their descriptions.

TABLE 4-7 Preconfigured Directives *(1 of 2)*

Directive Resource	Description
DOS standard	Used to back up and compress your Windows 95 and Windows 98 clients. By default, this resource has no directives. You can add directives as necessary.
DOS standard with compression	Used to back up and compress your Windows 95 and Windows 98 clients. Applies the `compressasm` ASM to all files.
NetWare standard	Used to back up your NetWare clients. By default, this resource has no directives. You can add directives as necessary.
NetWare standard with compression	Used to back up and compress your NetWare clients. Applies the `compressasm` ASM to all files.
NT standard	Used to back up your Windows NT and Windows 2000 clients. By default, this resource has no directives; you can add directives as necessary.

TABLE 4-7   Preconfigured Directives *(2 of 2)*

Directive Resource	Description
NT standard with compression	Used to back up and compress your Windows NT and Windows 2000 clients. It applies the `compressasm` ASM to all files.
UNIX standard	Contains a set of directives used to back up standard UNIX clients. Applies the following ASMs: • The `skip` ASM is applied to the `tmp_mnt` directory. • The `skip` ASM is applied to `core` files on the filesystem. • The `allow` save environment keyword is applied to the `/nsr` directory to ensure that local directives in `/nsr` and subsequent subdirectories are applied. • The `skip` ASM is applied to the `/tmp` directory. • The `swapasm` ASM is applied to the `/export/swap` directory. If swap files are located in a different directory, modify this directive to use the appropriate directory. • The `logasm` ASM is applied to the `/nsr/logs`, `/var`, `/usr/adm`, and `/usr/spool` directories. You may want to apply this ASM to other directories as well. • The `mailasm` ASM is applied to the `/usr/spool/mail` and `/usr/mail` directories. If mail files are located in different directories, modify these directives to use the appropriate locations.
UNIX standard with compression	Contains the same set of directives as the UNIX standard Directive resource, along with applying the `compressasm` ASM to all files.

▼ To Create a Directive Resource

To create a Directive resource:

1. In the Solstice Backup Administrator program, select Customize>Directives to open the Directives window.

2. Click Create.

3. In the Name attribute, enter a name for your new directive.

4. In the Directive attribute, enter one or more directives. For instructions on the syntax to use when entering directives, see "Directive Format" on page 183.

5. Click Apply.

The Directive resource is now available to be selected in the Directive attribute for Client resources.

▼ To Edit a Directive Resource

Note – You cannot change the name of a directive.

To edit a Directive resource:

1. In the Solstice Backup Administrator program, select Customize>Directives to open the Directives window.

2. From the Directives list box, select the appropriate directive.

3. Modify the directive as necessary and click Apply.

▼ To Delete a Directive Resource

Note – You cannot delete preconfigured directives or any directives currently applied to a Client resource.

To delete a Directive resource:

1. In the Solstice Backup Administrator program, select Customize>Directives to open the Directives window.

2. Select the appropriate directive in the Directives list box and click Delete.

   When prompted, click OK to confirm the deletion.

Adding or Changing the Solstice Backup Servers for this Client

The Solstice Backup software uses the contents of the `/nsr/res/servers` file to control who has client tasking rights (the right to request a program to be executed on another client). This tasking might be any of the following:

- Server performing an archive request
- Scheduled backup
- Migration
- Another client requesting a directed recover

During Solstice Backup software installation, you have the opportunity to add the names of Solstice Backup servers to this file. If you want to add additional servers at a later date, you must use a text editor to add the server name to the file.

If you want the client with the tasking rights to back up to other Solstice Backup servers, you must add the names of the additional Solstice Backup servers to this file. If you want other clients to be able to perform directed recovers to the client with the tasking rights, you must also add their names to the servers file.

Caution – If the servers file is empty, any Solstice Backup server has tasking rights, and this is a potential security concern.

▼ To Edit the servers File

To give clients tasking rights by adding their names to the *servers* file:

1. Start a text editor and open the /nsr/res/servers file (the default installation location for this file).

2. Enter one server name per line.

3. Save your changes and exit the text editor.

4. Restart the nsrexecd daemon.

Client Priority

The Client Priority attribute in the Client resource specifies the order in which participating clients are probed for the information needed to complete the save set worklist for that client. The Client Priority attribute can contain a value between 1 and 1,000. The lower the value, the higher the priority.

The client with the lowest value for the Client Priority attribute is placed at the top of the list to be contacted by the Solstice Backup server. If you do not specify a value in the Priority attribute, the contact order is random.

While the Client Priority attribute specifies the order of client contact, many variables affect the order in which clients complete their backups, including the following scenarios:

■ The backup operation on a client does not begin until the worklists for each of the save sets on the client are complete.

- The amount of work can vary greatly from one client to the next.
- If a client hangs and times out, it is put at the end of the list of clients to be contacted.

  To increase the number of times each client in a group is retried before the backup attempt is considered unsuccessful, change the value in the Client Retries attribute in the Group resource.

Solstice Backup Client Licensing

When a backup begins, a Solstice Backup client passes its attributes to the Solstice Backup server. The Solstice Backup server uses this information to verify that the client is allowed to back up to it. The following read-only attributes in the Client resource are used for enforcing client licensing:

- Client OS type
- Solstice Backup version
- CPUs
- Enabler in use

When you enable a Solstice Backup server, the enabler also allows you to use a specific number of clients of the same platform type. For example, a Solstice Backup server on Solaris is licensed for a specific number of Solaris or SunOS clients. If you want to connect clients of other platforms to that Solstice Backup server, you must purchase a Legato NetWorker ClientPak enabler for each of those platforms.

You can define more clients than the number of clients allowed by all enabler codes stored in the server. However, the Solstice Backup server enforces the number and type of licensed clients at backup time by verifying:

- The number of connected clients is less than or equal to the number of clients allowed by its enabler codes.
- The clients are of the types allowed by the enabler codes.

Sun no longer makes a distinction between different types of client connection licenses. Only *universal* licenses are sold. These licenses are valid for any computer, regardless of the computer's operating system.

The **nsrlic** Program

The `nsrlic` program generates reports about all licenses currently active on a Solstice Backup server. The program shows a breakdown of the types of workstations and servers that are in use. In addition, `nsrlic` has a `-v` option, which provides a list of all connected clients, as well as additional license information. For information about the `nsrlic` command, refer to the *Legato Command Reference Guide*.

The reports generated by the `nsrlic` program still separate client licenses into universal and workstation licenses.

- All computers with operating systems that could be a server (for example, Solaris, AIX or Windows NT server) are listed under universal.
- All others (for example, Windows 95, Windows 98, or Windows NT) are listed as workstations.

Workstation licenses will be "borrowed" from the universal licenses, unless the server has old workstation client connection licenses.

Note – Clients running Solstice Backup software earlier than release 5.0 do not have "self-ID" capabilities and are, therefore, not able to report their operating systems to the server at backup time. Since the server does not know what operating system the client is running, it will always use a workstation license, even if the client is on a Windows or UNIX server. However, such non-ID clients have no effect on which ClientPak licenses are required to allow Solstice Backup software to function; the appropriate licenses are still required.

Troubleshooting Client Licensing

In the following cases, Solstice Backup backups will be rejected because of licensing enforcement, and an error message is generated:

- Backups will be rejected if the client type is not allowed by the enabler code, or if a ClientPak enabler is not present. For example, the following error message might be sent for a Solaris server with an unlicensed Windows NT client:

```
RAP error, ClientPak for NT does not exist in the system.
```

- Backups will be rejected if no more licenses are available for that type of client. For example, you might receive an error message similar to this:

```
RAP error, Too many clients. Maximum is 2
```

Device and Media Management

This chapter describes device and media operations you can perform through the Solstice Backup server.

This chapter contains the following sections:

Note – The term "autochanger" refers to a variety of backup devices: autoloader, carousel, datawheel, jukebox, library, and near-line storage.

Pools

A *pool* is a specific collection of media to which the Solstice Backup server writes data. The Solstice Backup server uses pools to sort and store data. The configuration settings for each pool act as filters that tell the server which volumes should receive specific data. Pools are used in conjunction with label templates to keep track of which data is on which specific volume. For detailed information about label templates, see "Volume Labeling" on page 208.

This section covers the following topics:

- "How the Solstice Backup Server Uses Pools" on page 194
- "Solstice Backup Pool Types" on page 195
- "How the Solstice Backup Server Uses Pool Criteria to Sort Data" on page 197
- "Directing Client Indexes and Bootstrap to a Separate Pool" on page 198
- "Meeting the Criteria for More Than One Pool Configuration" on page 199
- "Not Meeting the Criteria for Any Pool" on page 200
- "Configuring Pools" on page 201
- "Using Storage Devices and Pool Configuration to Sort Data" on page 203
- "To Create a Pool" on page 204
- "To Edit a Pool" on page 206
- "To Delete a Pool" on page 207
- "To Create an Archive Pool" on page 207

How the Solstice Backup Server Uses Pools

The way you configure pools determines which volumes receive data. Each pool configuration contains criteria that the data must meet for the data to be written to associated volumes.

When a scheduled backup occurs, the Solstice Backup server tries to match the save set to a pool configuration. If the save set matches the criteria of a pool configuration, it is directed to a labeled volume from that pool. The server then checks if a correctly labeled volume for that pool is mounted on a storage device.

- If it is, the Solstice Backup server writes data to the volume.
- If it is not, the Solstice Backup server requests that such a volume be mounted and waits until an operator mounts the appropriate volume.

Solstice Backup Pool Types

The Solstice Backup software provides preconfigured pool types to keep different types of data separate. Solstice Backup software does not mix the following types of data on volumes within a pool:

■ Backup data

■ Archive data

■ Clone data

■ Migration data

Unless you specify other pools, all backup data is routed to the Default pool and all archive data is routed to the Archive pool. Cloned backup data is routed to the Default Clone pool, and cloned archive data is routed to the Archive Clone pool.

Preconfigured Pools

TABLE 5-1 lists the preconfigured pools that come with the Solstice Backup product.

TABLE 5-1   Preconfigured Pools

Type of Pool	Type of Data
Default	All backed up data when another pool is not specified.
Default Clone	Data from clones of the Default pool's volume.
Archive	Data from UNIX Archive clients (a UNIX archive is an archive that does not put entries in the client file index).
Archive Clone	Data from clones of the Archive pool's volume.
Full	Data from Full backups.
Migration	Migrated data.
Migration Clone	Data from clones of the migration pool's volume.
NonFull	Data from incremental and other non-Full backups.
Offsite	Data to be stored offsite.
PC Archive	Data from PC Archive clients. Unlike a UNIX archive, the PC archive puts entries in the client file index. (This is supported for Windows-based client only.)
PC Archive Clone	Data from cloned PC Archive clients.

Default Pool

If the save stream does *not* match the criteria for any custom (user-created) pools, the Solstice Backup server directs the save stream to the *Default* pool. Even if you are using your own customized pools or other preconfigured pools, always ensure that appropriate storage devices and media are available for your Default pool for those cases when the Solstice Backup server directs save streams there. If you do not have your Default pools properly configured for backup, you may receive an error message like the following in the Pending window:

```
media waiting (or critical): backup to pool 'Default' waiting for
1 writable backup tape or disk
```

Example: Using Pool Configurations to Sort Data

FIGURE 5-1 illustrates how the Solstice Backup server uses pool configurations to sort data. The save stream contains data from a full backup that was performed on client save sets in a group called *Accounting*. The Solstice Backup server looks for a pool configuration that matches the group *Accounting* and the level *Full*. When the server finds the matching pool configuration, it writes the data to a volume with a label from the corresponding *Accounting Full* pool of volumes mounted on one of the storage devices.
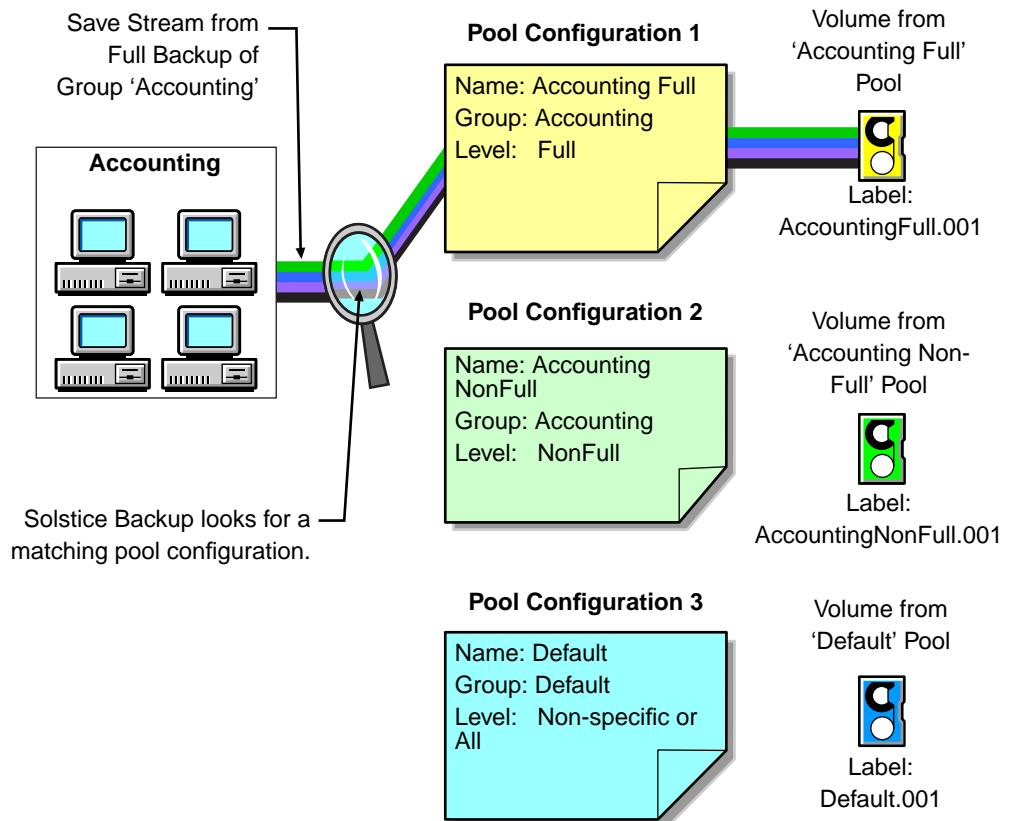
FIGURE 5-1   Using Pool Configurations to Sort Data

How the Solstice Backup Server Uses Pool Criteria to Sort Data

When you configure the Solstice Backup server, you can create additional pools and sort data by pool type and any combination of the following criteria:

- Group (backup group)
- Solstice Backup client
- Save sets (file or filesystems)
- Backup levels (full, levels [1–9], incremental, manual)

If you begin by entering a group name in the Group attribute, the pool is immediately restricted to accept only data associated with the named group. If you add a second group name to the Group attribute, the pool accepts data associated with either group, but no others. Entries for a single attribute function as "OR" clauses; that is, the pool accepts data from clients in either group.

Each of the four configuration criteria, however, functions with the others as an "AND" clause. That is, if you enter configuration criteria in both the Group attribute and Save Set attribute, only data that meets both the group criteria *and* the save set criteria is written to volumes from the specified pool.

You cannot create pools that share identical settings for the following:

■ Pool type
■ Group
■ Client
■ Save set
■ Level

If the settings for a new pool match the settings for an existing pool, you receive a warning message. Change the appropriate settings and reapply to save the Pool resource.

Directing Client Indexes and Bootstrap to a Separate Pool

You can use regular expression matching to direct the client indexes and bootstrap to a different pool than you send the backup data.

In the following example, the client file indexes are in index. To send the Solstice Backup server's bootstrap and all the client file indexes from this filesystem to the same pool, create a pool (in the Pool resource) with the following attributes:

TABLE 5-2   Example Attributes for Sending Bootstrap to a Separate Pool

Attribute	Value
Name	Index
Pool Type	Backup
Save Sets	bootstrap index

When the group's scheduled backup runs, the client save sets are written to a volume labeled for the appropriate save set pools, while the Solstice Backup server's bootstrap and `index` save sets are written to a separate volume labeled for the *Index* pool.

Meeting the Criteria for More Than One Pool Configuration

Depending on the pool configurations you create, you might have data that matches the criteria for more than one pool configuration. For example, if you configure one pool to accept data from a group called *Accounting* and you configure another pool to accept data from all full backups, the Solstice Backup server has to determine to which pool a full backup for the Accounting group is written. The server uses the following pool selection criteria:

1. Group (highest precedence)

2. Client

3. Save set

4. Level (lowest precedence)

When data matches the attributes for two pools, for example, Group and Level, the pool data is written to the pool specified in the Group attribute. For example, in the case where the data from the group matched the criteria for two different pools (one configured to accept data from the Accounting group and the other to accept data from all full backups) the data is routed to the pool that accepts data from the Accounting group.

TABLE 5-3 details the hierarchy that the Solstice Backup server uses to determine pool selection when a conflict arises. For example, the pool criteria for Group takes precedence over the pool criteria for Client, Save set, and Level. Data that meets the criteria for both pools is written to the pool associated with the Group. If data does not meet the criteria for any customized group, it is written to the Default group.

TABLE 5-3   Solstice Backup Hierarchy for Resolving Pool Conflicts *(1 of 2)*

Precedence	Group	Client	Save Set	Level
Highest	x	x	x	x
	x	x	x	
	x	x		x
	x	x		
	x		x	x

TABLE 5-3 Solstice Backup Hierarchy for Resolving Pool Conflicts *(2 of 2)*

Precedence	Group	Client	Save Set	Level
	x		x	
	x			x
	x			
		x	x	x
		x	x	
		x		x
		x		
			x	x
			x	
Lowest				x

Not Meeting the Criteria for Any Pool

When you use customized pool configurations to sort your data, you might inadvertently omit a client or save set. During a scheduled backup, if data does not meet the criteria for any customized pool configuration, the Solstice Backup server automatically sends the data to the Default pool. The server uses the Default pool to ensure that all data for clients in a backup group is backed up to a volume.

When the Solstice Backup server sends data to the Default pool, it looks for a labeled volume from the Default pool mounted on a storage device. If no Default pool volume is mounted on a storage device, the server requests the appropriate volume and waits until an operator mounts the volume. If the Solstice Backup server asks for a Default pool volume in the middle of a scheduled backup, but an operator is not present to mount it, the backup pauses until it has been mounted. If you have an operator available to monitor the backups, it is a good idea to keep a volume labeled for the Default pool close at hand in case this situation arises unexpectedly.

f you plan to use Solstice Backup software for unattended backups, run a test of the backup after making any configuration changes to ensure that all data is written to the appropriate volumes and to avoid an unexpected Solstice Backup request for a Default pool volume. For the procedure to test your scheduled backup, see "To Start a Group Immediately" on page 263.

Configuring Pools

The following sections provide information specific to configuring the following types of pools:

- Incremental backups
- Manual backups
- Clone data
- Archive data
- Migration data

Incremental Backups

If you want to create a separate pool for incremental backups, be aware that the Solstice Backup hierarchy of precedence affects the way the data is stored. If the Level attribute value is *incremental*, incremental data is routed to the associated pool but the corresponding changes to the client's file index are not. The Solstice Backup server saves all client file indexes at level 9 to speed the recovery operation, if one is needed.

If the client file indexes do not meet the criteria for the pool associated with the incremental backups, the Solstice Backup server:

- Matches the indexes to another pool (usually the Default pool).
- Looks for an appropriately labeled volume to write to.

If you need to recover all of your data, you need to use a large number of volumes. To store the client file indexes along with the incremental backup data and to speed the recovery operation, define the Level value in the Pool resource to accept both level 9 and incremental data.

You can use the Solstice Backup preconfigured NonFull pool settings to ensure that the client file indexes belong to the same pool as their incremental backups. When you keep the indexes in the same pool as their incremental backups, you reduce the number of volumes you need for a recovery.

Manual Backups

You can create a customized pool to receive data from a manual backup by specifying *manual* in the Level attribute. The Solstice Backup server, however, sorts data from a manual backup differently than data from a regularly scheduled backup. Because a manual backup is not performed as part of a scheduled backup group, the data is not associated with any group name. Thus, when you perform a manual backup in which only a single client's save set data is saved, the group normally

associated with that client's save set is not included as a criterion for pool assignment. As a consequence, data from a manual backup may be sent to a different pool than the pool in which data from this client's save set is stored during a regularly scheduled backup operation.

If you do not create a customized pool to receive data from manual backups, the Solstice Backup server uses the Default pool and looks for a mounted volume from the Default pool on which to write manually backed-up data. Because the Solstice Backup server tracks the volume location of all backup data, you do not need to worry about tracking which volume contains the manually backed-up data. If you need to recover the data, the server requests the correct volume.

When you perform a manual backup, the media database and resource files contained within the server bootstrap are not included in the backup. If you never perform regularly scheduled backups of the clients and server computers, the information vital to data recovery in the event of a disaster is not available. Refer to the *Disaster Recovery Guide* for further information on how the bootstrap is used during disaster recovery.

Clone Data

If you want to clone data, the Solstice Backup server requires a specific pool to receive the clone data and a minimum of two devices:

■ One to read the source volume

■ One to write the clone.

If you do not associate data to be cloned with a customized Clone pool, the Solstice Backup server automatically uses the Default Clone pool. You must mount an appropriately labeled volume on a separate storage device for the cloning process to proceed smoothly. For more information on the Solstice Backup cloning feature, see Chapter 9.

Archive Data

If you want to use Networker Archive to archive data, the Solstice Backup software requires a specific pool to receive the archived data. You can then store these volumes off site. If you do not associate data to be archived with a customized Archive pool, the Solstice Backup server automatically uses the preconfigured Archive pool. You must mount an appropriately labeled volume on a storage device for the archive process to proceed smoothly. For more information on the Solstice Backup archive feature, see Chapter 8.

Migration Data

Archive and migration data are both written in a different format than regular Solstice Backup save set data. Therefore, archive and migration data must each be written to different volumes. The client file indexes created during a PC archive, premigration, or migration operation are kept online and will not be written to the media until the next savegroup. By default, they are written to a volume from the Default pool every time the savegroup involving those clients is performed.

Using Storage Devices and Pool Configuration to Sort Data

You can sort data by configuring pools of volumes in conjunction with your storage devices. You can either use specific media to receive data, or designate a specific storage device to receive data from a designated pool.

Using Different Media

You can write data across several volumes of different media types (for example, magnetic disk and tapes) as long as the volumes mounted on the storage devices have the appropriate label associated with the pool. For more information about how a single save set can be written to more than one type of storage media, see "Client and Save Set Combination" on page 153.

Using a Specific Device for Backup Data

You can associate a pool with a specific storage device. For example, you may want your full backups written to optical disk for offsite storage. There are two ways to ensure that data goes to one specific storage device:

- Always keep a labeled volume associated with the appropriate pool mounted on the specific storage device.
- Associate a specific pool with the storage device in the Device resource. Then, in the Pool resource, associate the pool with the device in the Pool Configuration attribute list. All data will be written only to that device.

▼ To Create a Pool

To create a pool:

1. Create a label template for your pool. You also have the option of using a preconfigured label template. For a complete discussion of label template, see "How the Solstice Backup Server Uses Label Templates" on page 208.

Note – If you do not select a label template when creating a pool, the Solstice Backup server notifies you that it will create a label template for your pool. To have the Solstice Backup server create the label template, click OK.

---

2. From the Media Menu, select Pools.

3. Click Create.

4. Enter a name for the pool in the Name attribute. Because a pool is associated with a label template, you should use a name that matches the label template you created in Step 1. If you are using a preconfigured label template, you should name the pool in a way that clearly associates the pool with the corresponding label template.

5. From the Enabled attribute, select Yes.

6. From the Pool Type attribute, select the appropriate pool type.

7. In the Label Template attribute, select the matching label template.

8. Select one or more of the following: group, client, save set, or level. For information on the syntax of save set paths, see "Expression Matching of Save Sets to Pools" on page 205.

9. Select the devices that are permitted to mount volumes associated with this pool (optional).

10. In the Store Index Entries attribute, indicate whether client file index entries are to be generated for this pool.

11. In the Auto Media Verify attribute, select whether automated media verification will be performed while data is written to a volume associated with this pool.

12. In the Recycle to Other Pools attribute, select whether volumes marked as recyclable can be used by other pools.

13. In the Recycle from Other Pools attribute, select whether this pool can recycle volumes from other pools.

14. (Optional) Select a preference for volume type in the Volume Type Preference attribute.

The value you state is used as a selection factor when a request is made for a writeable volume.

15. Complete other attributes as necessary, and click Apply.

If any of these settings for a new pool match an existing pool, the following warning message appears:

```
Pool pool_name has overlapping selection criteria.
```

If you receive this message, review the pool configuration, change overlapping criteria (such as two pools configured for the same group or client and using the same backup levels), and click Apply.

Expression Matching of Save Sets to Pools

When you specify save sets to include in a pool, you can specify exact save set names, or you can use regular expression matching to send a group of save sets to a specific pool.

Regular expression matching affects how the save set path is written. Backslashes (common in Windows systems) and periods must be escaped with the backslash character to avoid being interpreted as part of the regular expression. For example, an entry reading:

```
/space/a.log
```

Would match /space/axlog, /space/ablog, /space/a.log, and so forth. To match only the file a.log, this would have to be written:

```
/space/a\.log
```

Similarly, the Windows directory C:\SomeDir\MyFiles would have to be written:

```
C:\\SomeDir\\MyFiles.
```

When using the Solstice Backup Administrator program, each save set in the Save Set attribute must be on a separate line. The following is an example of properly written save set entries:

```
/
/usr
/space/[jl]unk
c:\\Program Files\\bin
```

Refer to the *Legato Command Reference Guide* and the nsr_regexp and nsr_pool man pages for more information on these commands and expression matching.

Auto Media Verification

If the Auto Media Verify attribute in the Pool resource is enabled, the Solstice Backup server verifies data written to volumes from this pool by the following methods:

■ Data is verified by repositioning the volume to read a portion of the data previously written to the media. The data read is compared to the original data written. If the data read matches the data written, verification succeeds; otherwise it fails.

■ Media is verified whenever:

– A volume becomes full while saving and it becomes necessary to continue on to another volume.

– A volume goes idle because all save sets being written to the volume are complete.

When a volume fails verification, it is marked *full* so that the Solstice Backup server will not select that volume for future saves. The volume remains full until it is recycled or a user marks it *not full*. If a volume fails verification while the Solstice Backup server is attempting to switch volumes, all save sets writing to the volume are terminated.

▼ To Edit a Pool

You cannot change the name of a pool.

Note – The following preconfigured pools cannot be modified: Archive, Archive Clone, Default, Default Clone, Migration, Migration Clone, PC Archive, and PC Archive Clone.

To edit a pool:

1. From the Media Menu, select Pools to open the Pools window.

2. In the Pools attribute, select the pool you want to edit.

3. Edit the pool's attributes, and click Apply.

▼ To Delete a Pool

Note – Preconfigured pools cannot be deleted. You can delete other pools as long as there are no active volumes assigned to the pool.

To delete a pool:

1. From the Media Menu, select Pools to open the Pools window.

2. In the Pools attribute, select the pool you want to delete from the list displayed.

3. Click Delete.

4. When prompted, click OK to confirm the deletion.

5. Click Apply.

▼ To Create an Archive Pool

Note – An appropriately labeled volume must be mounted on a storage device for the archive process to proceed. For more information, see Chapter .

To create an Archive pool:

1. Create a new pool. For instructions, see "To Create a Pool" on page 204.

2. Set both the Pool Type to Archive and the Store Index Entries attributes to No. These two settings distinguish archive pools from backup pools.

As a result, when you schedule a new archive request, you may use the new Archive pool you created. If you choose to clone archived data, you can also create a new Archive Clone pool. The Solstice Backup server will write archive data only to an archive volume, and archive clone data only to an archive clone volume. For more information on pools, see "Pools" on page 194.

Volume Labeling

The Solstice Backup software labels (initializes) each storage volume with a unique internal label that corresponds to a pool. During backup and other operations, Solstice Backup software identifies the pool to which a volume belongs by its label. Solstice Backup software applies a label template to create a unique internal label for each volume.

The Solstice Backup server uses label templates and pool configuration settings to sort, store, and track data on media volumes. If you need to recover data, the server prompts you for the specific volume that contains the required data, identifying the required volume by the name given to it when it was labeled.

How the Solstice Backup Server Uses Label Templates

The Solstice Backup server writes a given set of data to a specific pool. For the server to recognize that a particular volume correctly belongs to a particular pool, the volume must have an internal identification label that associates it with the correct pool. The contents of the volume label follow rules defined in a specific label template that you create in the Label Templates resource. You then associate a label template with a specific pool in the Pool resource. If you do not associate data with a specific pool, the Solstice Backup server uses the preconfigured Default pool and corresponding Default label template.

FIGURE 5-2 illustrates how a pool configuration uses its associated label template to label a volume. You must configure a label template before you configure the associated pool for your custom template to be available in the Pool resource.
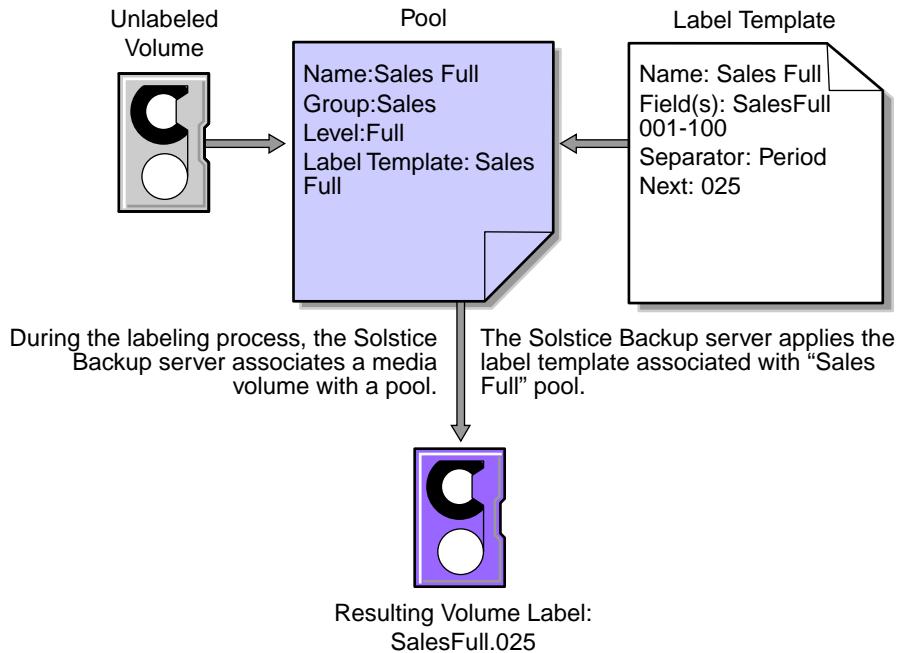
FIGURE 5-2 Labeling a Volume Using a Label Template

Preconfigured Label Templates

TABLE 5-4 lists the preconfigured label templates contained with the Solstice Backup server. These label templates correspond to the preconfigured pools.

TABLE 5-4 Preconfigured Label Templates *(1 of 2)*

Type of Label Template	Type of Data
Default	All backed-up data when another pool is not specified.
Default Clone	Data from clones of the Default pool's volume.
Archive	All archived data when another pool is not specified.
Archive Clone	Data from clones of the Archive pool's volume.
Full	Data from Full backups.
Migration	Migrated data.
Migration Clone	Data from clones of the Migration pool's volume.
NonFull	Data from incremental and other non-Full backups.

TABLE 5-4   Preconfigured Label Templates *(2 of 2)*

Type of Label Template	Type of Data
Offsite	Data to be stored offsite.
PC Archive	Data from PC Archive clients.
PC Archive Clone	Data from cloned PC Archive clients.
Two Sided	For two-sided volumes.

▼ To Complete Label Template Attributes

There are certain guidelines to keep in mind when completing the attributes for a
Label Template resource. TABLE 5-5 on page 210 describes how to complete the key
attributes for this resource. For more information on creating a label template, see
"To Create a Label Template" on page 214.

TABLE 5-5   Key Label Template Attributes *(1 of 2)*

Attribute	Guidelines
Name	Keep the label name consistent with the pool name, so that you can easily see how the data is organized. For example, you can create a label template called "AcctFull" for volumes that belong to a pool called "Accounting Full." The Solstice Backup software does not allow the following characters in label template names: / \ * [ ] ( ) $! ^ ' ; ' ~ < > & | { } In addition, you cannot use the following four characters, because they are used as separators in label templates: • colon (:) • dash (-) • period (.) • underscore (_)

TABLE 5-5 Key Label Template Attributes *(2 of 2)*

Attribute	Guidelines
Field	A label template is made up of one or more fields. Each field, or component, provides a layer of specificity to your organizational structure. You can have as many components as you want, but it is best to keep the template simple with few components. The label cannot exceed 64 characters. You can use four types of components: • Range of numbers (for example, *001-999*) • Range of lower case letters (for example, *aa-zz*) • Range of upper case letters (for example, *AA-ZZ*) • Character string (for example, *Accounting*) Each range includes a start value, a dash (-), and an end value. The start value and the end value must have the same number of characters. For example, use 01-99, *not* 1-99, or aaa-zzz, *not* aa-zzz. (This rule does not apply to a list of character strings or words; character strings should be separated by a blank space.) The order in which you enter each component of the Field attribute is important. The Solstice Backup server applies each component in a left-to-right order, starting with the first one entered. TABLE 5-6 on page 212 illustrates how label templates use components to create a number sequence for volume labels.
Separator	Choose which symbol you want to appear between component entries. You can use the period, dash, colon, or underscore to separate each component of the label template. If you do not select a separator, the label components do not have separators (for example, AA00aa), which makes the labels difficult to read.
Next	Choose the next sequence number to write on the label that the Solstice Backup server places on a volume (according to the template). • If you want to force a label to start the label scheme at a particular point, enter the start label value you want. The server continues to generate labels from that point on, according to the rules of the template. • If you want the Solstice Backup server to generate the first label for you, leave this attribute blank. When the Solstice Backup server recycles a storage volume, the volume label does not change as long as the volume remains in the same pool. That is, if a storage volume labeled "Dev.006" is recycled, it retains the volume label "Dev.006" and does not receive a new label with the next sequence number.

TABLE 5-6 lists examples of number sequences for volume labels.

TABLE 5-6　　Examples of Number Sequences for Volume Labels

Type of Components	Fields	Number Sequence Result	Total Number of Labels
Range of numbers	001-100	001, 002, 003,...100	100
• Character string • Range of numbers	SalesFull 001-100	SalesFull.001,...SalesFull.100	100
• Range of lower case letters • Range of numbers	aa-zz 00-99	aa.00,...aa.99, ab.00,...ab.99, ac.00,...ac.99, : az.00...az.99, ba.00,...ba.99 : zz.00,...zz.99	67,600 (26^2 times 10^2)

Your label template should allow for expansion of your backup media storage
system. For example, it is better to create a template for 100 tapes and not use all of
them, than it is to create a template for only 10 tapes and run out of labels. When the
Solstice Backup server reaches the end of the template number sequence, you will be
warned that you have reached the maximum tapes allowed for this label template.
Modify the label template to allow for more tapes.

Label Template Naming Strategies

The Solstice Backup server is packaged with preconfigured label templates that
correspond to the preconfigured pools. If you choose to create your own templates,
you can include as many components in the Fields attribute as necessary to suit your
organizational structure. However, it is a good idea to keep the template simple with
few components.

For example, if you create a label template for your accounting department, you can customize your label template in several ways, depending on the size of your storage system and media device capabilities. TABLE 5-7 illustrates several ways you can use components to organize your labels.

TABLE 5-7   Using Label Template Components

Type of Organizational Structure	Components	Separator	Resulting Volume Labels
Sequential	AcctFull 001-100	period	AcctFull.001 (100 total labels)
Storage oriented (for example, 3 storage racks with 5 shelves each, each shelf holding 100 tapes)	1-3 1-5 001-100	dash	1-1-001 This label is for the first tape in rack 1 on shelf 1. (1,500 total labels)
Two-sided media (for example, optical devices)	AcctFull 000-999 a-b	underscore	AcctFull_000_a (side 1) AcctFull_000_b (side 2) (2,000 total labels)

Labeling Tips

Naming schemes vary from site to site. One suggestion is to name the volumes with the name of the Solstice Backup server followed by a three-digit number, for example:

```
jupiter.001
```

No matter what naming scheme you choose to use, keep in mind that the simpler a convention is, the easier it will be understood by operators and administrators.

When you label volumes, each volume should have a physical (adhesive) label attached to it. Since the Solstice Backup server keeps track of the backups and which volumes they are on, you can name the volumes with any convenient name. For example, you can label your volumes *1, 2, 3,* or *Monday.1, Tuesday.1, Wednesday.1.* You can assign a volume any name as long as each one is unique.

The adhesive label on the volume should match the name generated by the Solstice Backup software. For example, if you physically label a volume *mars.1*, its Solstice Backup name should also be *mars.1*.

▼ To Create a Label Template

For information on the template attributes, refer to TABLE 5-5 on page 210.

To create a label template:

1. From the Customize menu, select Label Templates to open the Label Templates window.

2. An empty Name attribute appears. Click Create.

3. In the Name attribute, enter the name of the label template.

4. Enter the components for your label template in the Fields attribute. Press [Return] after entering each component.

5. In the Separator attribute, select a separator symbol. If no symbol is selected, the components will have no separators (for example, AA00aa).

6. For the Next attribute, enter the next label the Solstice Backup server will use to label the volume according to the template.

7. Click Apply.

▼ To Edit a Label Template

You cannot change the name of a label template. However, if you want to change an individual label name, delete the existing name in the Next text box, and enter a new name.

To edit a label template:

1. From the Customize menu, select Label Templates to open the Label Templates window.

2. In the Label Templates attribute, select the label template you want to edit.

3. Edit the attribute values of the label template, and click Apply.

▼ To Delete a Label Template

You cannot delete the preconfigured label templates or a label template currently in use.

To delete a label template:

1. From the Customize menu, select Label Templates to open the Label Templates window.

2. In the Label Templates attribute, from the list displayed, select the label template you want to delete.

3. Click Delete.

4. When prompted, click OK to confirm the deletion.

5. Click Apply.

━━━━━━

Devices

A device is a drive that reads and writes data to storage volumes or a disk file during backup, recovery, and other operations. You may add additional devices to the Solstice Backup server since the software supports multiple devices in a series. That is, it will use a tape in one device until it is full and then use the tape in the second device. If you add or remove a device from the Solstice Backup server, the server will automatically recognize and display the change when it is restarted.

You must configure each device before the Solstice Backup server can recognize and use it. The Device resource contains the attributes for each device that you must set. The instructions for configuring your devices differ depending on whether the device is stand-alone or is contained in an autochanger or silo.

Device Support by Solstice Backup Edition

Each edition of Solstice Backup software has a different maximum number of devices a server can administer. Each enabled storage node connected to the Solstice Backup server can increase the device maximum. The maximum number of devices for any combination of Solstice Backup server and storage nodes is 256. Device levels for other Solstice Backup platforms and releases can vary, so refer to the documentation of that particular server.

TABLE 5-8 displays the maximum number of devices allowed for this Solstice Backup software release, including the maximum for each enabled storage node.

TABLE 5-8   Maximum Devices

Solstice Backup Product	Server	Each Storage Node
Workgroup Edition	4	Not Applicable
Network Edition	16	16
Power Edition	32	32

If your server is configured to support more than one device, it will use the volume in a device until it is full, then use the volume in another device, and so on.

Use concurrent devices to spread backups evenly across your storage devices, and use the optional Networker Autochanger Software Module for unattended backups to an autochanger.

For a list of the currently supported storage devices, refer to the *Sun Compatibility Guide* on the Sun web site at *www.sun.com*.

▼ To View Supported Devices and Media Types

The Solstice Backup server supports a wide range a storage devices and media types, including tape, optical, and file type devices. The storage devices and supported media types can sometimes vary.

To view the media types available to your particular Solstice Backup server:

1. From the Media menu, select Devices to open the Devices window.

2. In the Media Type attribute, display the list of supported devices. This list displays abbreviations for all the device types supported in that Solstice Backup version.

3. From the Help menu, select On Window to view more detailed descriptions of each device type.

Filesystem devices are configured differently from traditional devices and are covered separately in this manual. For information on file device types, see "Filesystem Devices" on page 232.

Configuration of the Sun StorEdge Network Foundation HBA/Driver

The StorEdge™ Network Foundation HBA requires a special configuration to bind the device World Wide Port Names that are used by the Legato `lus` driver.

The `/etc/LGTOuscsi/lus_add_fp_devs` script adds the proper entries to the `lus.conf` files. However, before you run the script, you must ensure that all devices are powered on and connected to the HBA(s).

Note – The `/etc/LGTOuscsi/lus_add_fp_devs` script should only be used to configure libraries over fibre channel.

The following are examples of additional entries for the `/usr/kernel/drv/lus.conf` script:

```
name="lus" parent="fp" target=0 lun=0
fc-port-wwn="22000090a50001c8";

name="lus" parent="fp" target=2 lun=0
fc-port-wwn="22000090a50001c8";
```

Note – Each time you add a new device, or remove an old device, you must re-create and reinstall the new `lus` configuration. For information explaining how to add or delete a device, see "To Add a Device" on page 220.

Caution – When you run `/etc/LGTOuscsi/lus_add_fp_devs` script, you should see the World Wide Port Names displayed in the script output. If none are displayed, this might indicate your fiber channel is incorrectly configured and that you should check your configuration. Use the Sun `luxadm` program to check the fiber channel connection.

Sample `/etc/LGTOuscsi/lus_add_fp_devs` Script Output

The following is an example of `/etc/LGTOuscsi/lus_add_fp_devs` script output.

```
# /etc/LGTOuscsi/lus_add_fp_devs
Updating /usr/kernel/drv/lus.conf

Found StorEdge HBA device:
/devices/pci@1f,0/pci@1/pci@1/pci@f/SUNW,qlc@4/fp@0,0:devctl

Found StorEdge HBA device:
/devices/pci@1f,0/pci@1/pci@1/pci@f/SUNW,qlc@5/fp@0,0:devctl

Mapping tape devices starting at target 0

   Adding:
     port WWN: 500104f000428e48 as SCSI target 0
     port WWN: 500104f000428e49 as SCSI target 1
     port WWN: 500104f000428f44 as SCSI target 2
     port WWN: 500104f000428f45 as SCSI target 3

Mapping tape devices starting at target 4

   Adding:
     port WWN: 500104f000428e48 as SCSI target 4
     port WWN: 500104f000428e49 as SCSI target 5
     port WWN: 500104f000428f44 as SCSI target 6
     port WWN: 500104f000428f45 as SCSI target 7

Adding medium changer devices starting at target 100

   Adding:
     port WWN: 100000e002000000 as SCSI target 100
     port WWN: 100000e00201105a as SCSI target 101
Adding medium changer devices starting at target 102
```

```
    Adding:
      port WWN: 100000e002000000 as SCSI target 102
      port WWN: 100000e00201105a as SCSI target 103

The changes made to lus.conf will not take effect until lus is
reloaded.

If you are sure that no process is currently using lus to control
a tape library, then it may be safely unloaded.
```

Note – If *lus* is not currently loaded, proceed to the section in the script stating *lus* is not currently loaded.

```
Would you like to unload lus? (y/n)  y
```

Note – If you select no, the following output is displayed and the program exits:

```
lus is currently loaded! It must be unloaded and reloaded for any
changes to lus.conf to take effect
```

Note – If *lus* is not currently loaded, complete the following prompt:

```
lus is not currently loaded

Would you like to run inquire?
(this will cause lus to be loaded using the newly modified
lus.conf) (y/n) y
```

```
scsidev@0.0.0:FUJITSU MAJ3182M SUN18G 0804|Disk,
/dev/rdsk/c0t0d0s2
scsidev@2.0.0:STK 9840 1.28|Tape, /dev/rmt/0cbn
scsidev@2.1.0:STK 9840 1.28|Tape, /dev/rmt/2cbn
scsidev@2.100.0:ADIC Scalar DLT 448  0114|Autochanger (Jukebox)
scsidev@2.101.0:ADIC Scalar DLT 448  0114|Autochanger (Jukebox)
scsidev@3.0.0:STK 9840 1.28|Tape/dev/rmt/6cbn
scsidev@3.1.0:STK 9840 1.28|Tape, /dev/rmt/8cbn
```

Nonrewinding Tape Device Usage

If you use tape drives as your storage devices, you must use nonrewinding devices.
The Solstice Backup server writes a filemark on the volume at the end of each
backup and then appends data onto the volume based on the position of the
filemark. If the device rewinds the media, the filemark position will be lost and
previously written data is overwritten by the next backup. The pathnames for these
devices must follow the Berkeley Storage Device (BSD) semantic rules, for example,
/dev/rmt/0mbn. The "b" in the pathname satisfies the BSD semantics requirement.

▼ To Add a Device

You need to create a new Device resource for each stand-alone device or
autochanger device. When you elect to create a Device resource, some of its
attributes have default values. You can use these values if they apply to your storage
device, or change them to a different value.

Caution – The Solstice Backup server automatically creates a 8mm tape device
during its installation. If this device is incorrect, you will have to delete the existing
Device resource and create a new one with the correct media type. The Media Type
attribute of an existing Device resource cannot be changed. For information on
deleting devices, see "To Delete a Device" on page 222.

To add a device from the Solstice Backup Administrator Program:

1. From the Media menu, select Devices to open the Devices window.

2. Click Create.

3. In the Name attribute, enter the device pathname.

   If you use a file device, you must enter it as a directory path rather than only a filename.

4. From the Media Type attribute, select the appropriate media type.

   If you want to back up to a filesystem device configured on your disk, see "Filesystem Devices" on page 232 for specific instructions and restrictions. For more information on the logical selection, see "Logical Volume Backup" on page 233.

5. For the Read-Only attribute, select Yes to have the device read volumes only (instead of writing the data to volumes).

6. In the Target Sessions attribute, enter or select a value. For more information on how this attribute works, see "Target Sessions" on page 95.

7. In the Cleaning Required attribute, select Yes to enable the cleaning requirement feature.

   The cleaning device options are available only if you are backing up to an autochanger. This attribute is valid only for autochanger tape devices.

8. In the Cleaning Interval attribute, enter a date to set the period of time between cleanings. Note that the Cleaning Required attribute must be also enabled. Specify the time in days, weeks, or months. For example:

   ■ Four days

   ■ Two weeks

   ■ Month (same as one month)

   This attribute is only valid for autochanger tape devices.

9. In the Date Last Cleaned attribute, enter the date the device was last cleaned. The Solstice Backup server determines the next scheduled cleaning time using this date in conjunction with the value you entered in the Cleaning Interval attribute. This attribute displays the date of subsequent device cleanings.

   This attribute is only valid for autochanger tape devices.

10. In the Auto Media Management attribute, select Yes to enable this feature. For information on this feature, see "Auto Media Management" on page 235.

11. Complete other attributes as necessary, and click Apply.

▼ To Edit a Device

To edit a device from Solstice Backup Administrator:

1. From the Media menu, Select Devices to open the Devices window.

2. In the Devices attribute, select the device you want to modify from the list displayed.

3. Edit the attributes, and click Apply.

▼ To Delete a Device

Note – You cannot delete devices that have volumes mounted.

To delete a device from Solstice Backup Administrator:

1. From the Media menu, select Devices to open the Devices window.

2. In the Devices attribute, select the device you want to delete.

3. Click Delete.

4. When prompted, click OK to confirm the deletion.

5. Click Apply.

Device Parameter Environmental Variables

A number of environmental variables can be used to adjust several Solstice Backup device parameters:

- **NSR_DEV_BLOCK_SIZE_***DEVICE_NAME*
- **NSR_DEV_TAPE_FILE_SIZE_***DEVICE_NAME*
- **NSR_DEV_DEFAULT_CAPACITY_***DEVICE_NAME*
- **NSR_DEV_LOAD_TIME_***DEVICE_NAME*
- **NSR_DEV_LOAD_POLL_INTERVAL_***DEVICE_NAME*
- **NSR_DEV_LOAD_TRY_LIMIT_***DEVICE_NAME*

Where *DEVICE_NAME* is the backup device type available to your Solstice Backup server. Note that the syntax for the device name must be all upper case, with underscores (_) replacing blank spaces. Therefore, a device displayed in the Solstice Backup software as "4mm 20gb" would be listed as:

4MM_20GB

Environment variables should be set in the Solstice Backup startup script (/etc/init.d/networker) before the lines:

```
(echo 'starting Solstice Backup daemons:') > /dev/console
```

Enter the environment variables using the format:

```
ENV_VAR_NAME = value
export ENV_VAR_NAME
```

NSR_DEV_BLOCK_SIZE_*DEVICE_NAME*

NSR_DEV_BLOCK_SIZE_*DEVICE_NAME*_ is organized in units of KB. The value set must be a multiple of 32, with a minimum value of 32. Maximums are determined by platform, SCSI driver, and device. For example:

```
NSR_DEV_BLOCK_SIZE_4MM_20GB=64
```

NSR_DEV_TAPE_FILE_SIZE_*DEVICE_NAME*

NSR_DEV_TAPE_FILE_SIZE_*DEVICE_NAME*_ is organized in units of **NSR_DEV_BLOCK_SIZE**_*DEVICE_NAME*_ and is the number of blocks written between filemarks. These filemarks are used to locate a particular spot on the tape during recovery, and more filemarks will generally lead to faster positioning. For example:

```
NSR_DEV_TAPE_FILE_SIZE_TZ89=512
```

The Solstice Backup software writes a filemark by closing and reopening the tape device; requiring a second or two. If this value is too small, throughput may be slowed.

NSR_DEV_DEFAULT_CAPACITY_*DEVICE_NAME*

NSR_DEV_DEFAULT_CAPACITY_*DEVICE_NAME*_ is the size of the particular tape used to base the percent full calculation. This variable value has no effect on the actual tape capacity. Any integer value is allowed, with a *KB*, *MB* or *GB* designation

to indicate a range of values. Any value less than 200 MB will be overridden by the normal default capacity. There is no obvious maximum, with the only practical limitation being the actual storage size. For example:

```
NSR_DEV_DEFAULT_CAPACITY_DTL7000=12GB
```

NSR_DEV_LOAD_TIME_*DEVICE_NAME*

NSR_DEV_LOAD_TIME_*DEVICE_NAME* is the number of seconds that nsrmmd polls and waits for a drive to become ready after the library inserts a tape into the device. **NSR_DEV_LOAD_POLL_INTERVAL_***DEVICE_NAME* is used to set the number of seconds nsrmmd waits between polls during load time.

If the value of **NSR_DEV_LOAD_TIME_***DEVICE_NAME* is too short, there may be unnecessary load failures. If it is too long, then labeling new tapes will take longer than necessary. The minimum allowable value is 10 seconds, the maximum value is 600 seconds. For example:

```
NSR_DEV_LOAD_TIME_DTL8000=300
```

NSR_DEV_LOAD_POLL_INTERVAL_*DEVICE_NAME*

NSR_DEV_LOAD_POLL_INTERVAL_*DEVICE_NAME* is the number of seconds that **nsrmmd** waits between polls during load time. The minimum allowable value is 1 second, the maximum value is 30 seconds. For example:

```
NSR_DEV_LOAD_POLL_INTERVAL_DLT=10
```

NSR_DEV_LOAD_TRY_LIMIT_*DEVICE_NAME*

NSR_DEV_LOAD_TRY_LIMIT_*DEVICE_NAME* is the number of times that **nsrmmd** will attempt to open a drive. nsrmmd will poll the drive until the limit set in **NSR_DEV_LOAD_TIME_***DEVICE_NAME* has been reached. After the limit has been reached, it will retry until the **NSR_DEV_LOAD_TRY_LIMIT_***DEVICE_NAME* has been reached. The default value and minimum allowable value is 2, the maximum value is 120.

```
NSR_DEV_LOAD_TRY_LIMIT_DLT=4
```

Dynamic Drive Sharing

Dynamic Drive Sharing (DDS) is a feature that provides Solstice Backup software with the ability to recognize shared drives. DDS enables Solstice Backup software to skip the shared drives that are in use and route the backups or recoveries to other available shared drives.

This section contains the following topics:

- "Introduction to DDS" on page 225
- "DDS Configuration" on page 226
- "Block Size Compatibility" on page 227
- "Enabling DDS within NDMP" on page 228
- "DDS Attributes in the Devices Dialog Box" on page 229
- "To Create a New Device and Enable DDS" on page 230
- "DDS Licensing" on page 231

For information explaining how to enable DDS when configuring a device with the jbconfig program, see "To Enable DDS When Configuring an Autochanger" on page 374.

Introduction to DDS

DDS controls application requests for media and allows the Solstice Backup server and all storage nodes to access and share all attached devices.

A system administrator can enable DDS when creating a device through the Solstice Backup Administrator program or when configuring a device with the jbconfig **program.**

Two terms central to the use of DDS are *drive* and *device*. Within the context of DDS, these terms are defined as follows:

- Drive — The physical backup object, such as a tape drive, disk, or file.
- Device — The access path to the physical drive.

Note – Currently, DDS is only supported in a SAN fiber channel environment and not a direct-connect SCSI environment.

Benefits of DDS

Enabling DDS on your system provides you with the following benefits:

- Reduces storage costs — A single tape drive can be shared out among several storage nodes. In fact, since the Solstice Backup software uses the same open tape format for UNIX, Windows NT, NetWare and Linux, you can actually share the same tape between different platforms (assuming that respective save sets belong to the same pool).

- Reduces LAN traffic — You can reconfigure clients as SAN storage nodes which can send save sets over the SAN to shared drives.

- Provides fault tolerance — Within a SAN environment you can configure your hardware to eliminate a single point of failure.

- Provides configuration over a greater distance — Enables you to configure a system over a greater distance than with SCSI connections.

DDS Configuration

> ⚠️ **Caution –** DDS licenses must be installed before the `jbconfig` command is used to configure shared devices. See the Solstice Backup Roadmap for information about licensing DDS.

FIGURE 5-3 on page 227 illustrates the DDS process and potential configurations for sharing drives. This basic configuration consists of a server, two storage nodes, and an autochanger with two tape drives.

In this figure, storage nodes sn_1 and sn_2 are attached to the autochanger. Each node, on its own, has access to drive_1 and drive_2. With DDS enabled, both nodes have access to both drives and—more importantly—can recognize when a shared drive is in use.

For example, sn_1 sends a backup to drive_1 and subsequently, sn_2 starts a backup to drive_1. If the backup from sn_2 can use the same tape in drive_1, it will. If not, the backup started by sn_2 will be sent to the next available drive, which is drive_2.

Under such a configuration, two DDS licenses are required, one for each storage node.
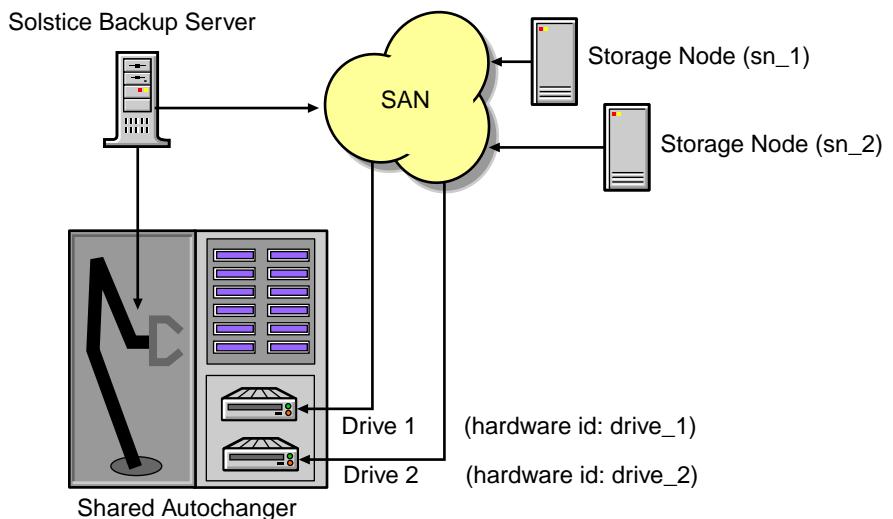
FIGURE 5-3   Dynamic Drive Sharing

Block Size Compatibility

With DDS enabled, you can share a drive between storage nodes on different platforms, such as UNIX and Windows. For Solstice Backup software operations, such as backups and recoveries, to take place successfully, you need to ensure that the block size is compatible between different platforms and/or hardware.

To ensure compatibility, make sure one of the following conditions is met:

- The various storage nodes sharing a drive support the same block sizes.
- When a tape is labeled on a drive, it is labeled with the block size which is the smallest of all block sizes supported by all the storage nodes sharing the drive.

For information on how to set block sizes for individual drives or tapes on different platforms, see "Setting Compatible Block Size for UNIX and Microsoft Windows" on page 386.

Block Size Incompatibility between UNIX and Windows

Incompatible block size settings between UNIX and Windows storage nodes may result in any one of the following error scenarios:

- A backup taken on a UNIX node may not be recoverable on a Windows node because the Windows node does not support large block sizes.

- A UNIX process labels and saves data to a tape and leaves the tape mounted. A Windows process subsequently attempts to verify the label on this tape and fails because the label verification is done by reading a header from the data portion.

- A tape on a UNIX node is labelled with a large block size. The backup is started on a Windows node and the Windows node attempts to write the backup using the default block size. Internally, the backup on Windows NT is written by breaking down the big buffer of data into smaller segments of writable blocksizes.

  Attempting to recover a specific file on Windows NT in this situation fails because of positioning errors on the tape.

Enabling DDS within NDMP

The following sections explain the requirements for successfully enabling DDS with NDMP.

DDS on NDMP Nodes in a SAN Environment

You can share drives between NDMP nodes in a SAN environment. However, you must ensure the NDMP nodes are homogeneous.

For example, you can enable DDS in the following configurations:

- EMC Celerra to EMC Celerra
- EMC IP4700 to EMC IP4700
- NetApp to NetApp (any NetApp nodes that Network Appliance supports within a zone)

The current NDMP implementation does not allow you to share drives between non-homogeneous NDMP nodes. However, there is no inherent limitation within DDS preventing this.

FIGURE 5-4 on page 229 illustrates a basic DDS configuration with NDMP.
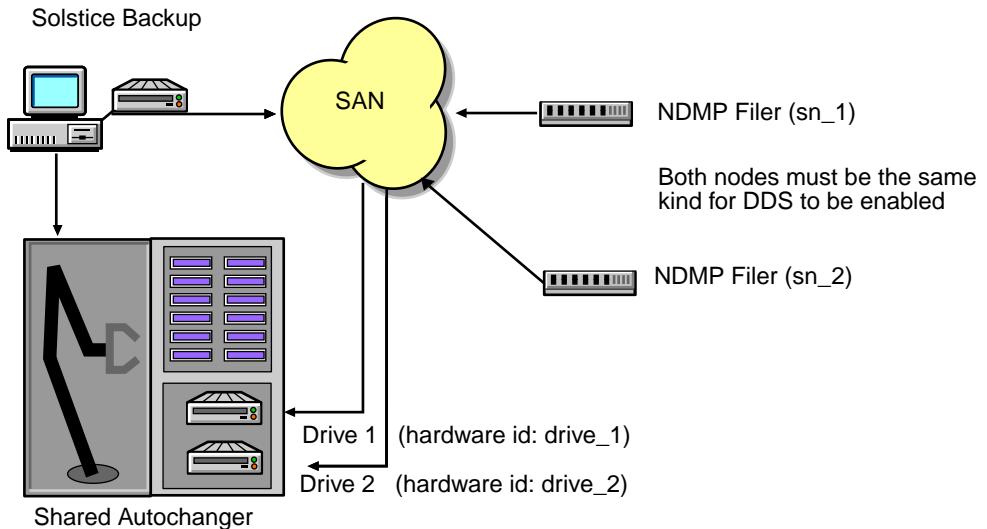
FIGURE 5-4 DDS with NDMP

DDS Attributes in the Devices Dialog Box

The Devices dialog box includes two attributes used in configuring DDS:

■ Hardware ID

■ Shared Devices

Hardware ID Attribute

The hardware ID tracks drives being shared by multiple hosts. Device instances sharing the same physical drive across multiple hosts have the same hardware ID. You can update the hardware ID during the device configuration process. Any character string can be used as a value.

If you are enabling DDS through the jbconfig program, the program prompts you for a hardware ID during the configuration process.

Shared Devices Attribute

The Shared Devices attribute features values which can be used to manipulate all shared instances of a drive at the same time (see TABLE 5-9). This attribute allows you to enable or disable all devices sharing the same hardware ID with a single action.

TABLE 5-9   Shared Devices Attributes

Value	Description
Enable All	When selected, enables all devices using the same hardware ID.
Disable All	When selected, disables all the devices using the same hardware ID.
Done	This is the default setting. After the server has enabled or disabled all devices with the same hardware ID, the attribute is reset to done.

The Shared Devices attribute is not reflected in the jbconfig program.

▼ To Create a New Device and Enable DDS

You can enable the DDS feature when you create a device. If you assign a non-unique hardware ID to the device in Step 4, you also enable DDS—when another device has the same hardware ID assigned to it.

Device instances sharing the same physical drive across multiple hosts should have the same hardware ID. This indicates to Solstice Backup software that these multiple devices are the same physical drive.

To create a device and enable DDS:

1. From the Media menu, select Devices to open the Devices window.

2. Click Create.

3. Enter the device pathname into the Name attribute and complete the remaining fields as required for your system. For information explaining how to create a new device, see "To Add a Device" on page 220.

4. Assign a non-unique hardware ID to the device. The hardware ID accepts any character string as a value.

5. Complete any other attributes as necessary and click Apply.

▼ To Stop Sharing a Device through the Devices Window

To stop sharing a drive that has been assigned a Hardware ID, remove the value from the hardware ID attribute in the Devices window.

⚠ **Caution –** When you stop sharing a DDS-configured drive, this does not mean that the drive is disabled. This means that Solstice Backup software is no longer aware that a specified drive is shared.

Consequently, so as not to impede Solstice Backup operations when you stop sharing a device, you should remove the other device instances.

▼ To Change the Idle Device Timeout Attribute

You might notice that a tape remains mounted in its drive for 10 to 15 minutes after a backup has completed. Other requests for the drive from another device path will wait during this period.

You can modify this timeout value by adjusting the Idle Device Timeout attribute. To change this attribute:

1. In the Solstice Backup Administrator program, select Media>Jukeboxes to open the Jukeboxes window.

2. From the View menu, select Details.

3. Scroll down to Idle Device Timeout.

4. Select a time that is appropriate for your system configuration.

   The default Idle Device Timeout value is 10 minutes.

Note – Selecting zero means that the device never times out. Consequently, you would have to eject the tape manually.

DDS Licensing

To enable DDS, you need one license per drive that is to be shared. Once a drive is licensed as shared, any number of storage nodes can share it. DDS licensing is independent of autochanger and storage node licensing.

Filesystem Devices

One Solstice Backup data storage option is to save data to a computer's local disk, rather than use a tape or optical device to write the backed-up data to storage media. This method is known as *filesystem device* storage. Storing data to the filesystem greatly reduces the time it takes to both save and recover data.

You designate the Solstice Backup server to back up data to a filesystem device by creating a new Device resource for your Solstice Backup server the same way as for any other storage device. For instructions on creating device resources, see "To Add a Device" on page 220.

You can have more than one file type device backing up a filesystem, as long as they have unique names. Each file type device counts as one storage device for the Solstice Backup license, so be sure you do not exceed your device limit. Additionally, if you have multiple file type devices backing up to the same disk, you should partition the disk and limit one device per partition.

Use with Staging

Saving data to a file type device greatly increases the speed of backups and recoveries, but also uses a large amount of disk space. When backing up data to file type devices you may want to have a staging policy in effect to move the data off the disk as soon as possible to make room for the next backup. For more information on creating staging policies and using staging in conjunction with file type devices, See "Save Set Staging" on page 251.

▼ To Configure Filesystem Devices

Creating a filesystem device requires you to create a directory in the disk filesystem that the Solstice Backup server will recognize as the device name (and the destination for the data), for example:

```
/export/backup.data
```

Note – Do not use the `tmp` directory for filesystem devices because the data might be overwritten.

■ You might have problems using file type devices together with other types of devices in the same pool.

To configure a filesystem device:

1. Create one directory for each disk (or partition) being used for file type device storage.

2. In the Solstice Backup Administrator, select Media>Devices to open the Devices window.

3. Click Create.

4. In the Name attribute, enter the complete pathway for the directory.

5. For the Media Type attribute, select File from to indicate that the device is a filesystem device.

6. For the Target Sessions attribute, enter or select a value.

7. Leave the options for cleaning at their default disabled settings so that automatic cleaning is not invoked.

8. Do not enable the Auto Media Management feature for filesystem devices.

9. Complete other attributes as necessary, and click Apply.

You label and mount the file device in the same manner as for other offline media used for Solstice Backup backups. Use of a file as a device facilitates save set staging. For more information, see "Save Set Staging" on page 251.

Logical Volume Backup

A *logical volume* is a type of primary (disk) storage on a client computer that can span several physical disk volumes. The logical volume has its own device address, and it is treated similarly to a disk partition by the filesystem. When the Solstice Backup server backs up data from clients, it has to determine how many save sessions to allocate to each client for best performance. To avoid contention, there should not be more than one backup operation running per physical disk. Different sessions are allocated across different physical disks for this reason.

To determine how many save sessions to allocate, the Solstice Backup server probes (queries) the clients in a backup group using the savefs -p command. This command determines:

■ What data to back up.

■ Where the data is physically located.

■ Whether there are logical volumes.

The Solstice Backup server stores this information in two variables, *disk-number* and *maximum-sessions*, according to the following rules:

■ When the group of volumes or disks that contain logical volumes is not part of the device path, all logical volumes on the client computer are assigned to the same disk-number, and maximum-sessions is set to the number of logical volumes on the client computer.

■ When the group of volumes or disks that contain logical volumes is part of the device path, all logical volumes within the volume group are assigned to the same disk-number, and maximum-sessions is set to the number of logical volumes within the volume group.

The server uses the output from the `savefs` probe to allocate its save sessions (up to the maximum server parallelism) across the clients in the backup group.

It allocates the save sessions as follows:

1. The server allocates one save session per client in the backup group.

2. If there are still save sessions available, it allocates one save session per physical disk on each client.

3. If there are still save sessions available, it allocates save sessions to each *disk-number* value, up to the limits in *maximum-sessions* for each client and client parallelism.

Media Handling Error Resolution

A damaged tape can produce media handling errors. Trying to mount and read a damaged tape can result in a continuous loop of failed actions. The device can continuously try to mount the tape, replace it in the slot, and then retry the action again with the same result.

▼ To Limit Repeated Device Errors

To minimize the number of instances that this type of loop can occur:

1. From the Media menu, select Devices to open the Devices window.

2. From the View menu, select Details to display the hidden attributes.

3. In the Max Consecutive Errors attribute, lower the value (set to 20 by default), and click Apply.

When the number of errors meets the value set in this attribute, the device stops attempting the activity and becomes disabled.

4. Once the host running the Solstice Backup server is idle, remove any media from the disabled device and ensure that the device is in good working order.

5. With the Devices window still open, re-enable the device.

6. From the list in the Devices attribute, select the disabled device.

7. In the Enabled attribute, select Yes.

8. In the Write Enabled attribute, select *Yes,* unless you want the volume to be read-only.

9. Click Apply.

Note that if your disabled device is part of an autochanger, you might need to reset the device by opening a terminal window, changing your path to the directory containing the Solstice Backup executables, and entering the following command:

```
# nsrjb -HE
```

Auto Media Management

The auto media management feature gives the Solstice Backup server automatic control over media loaded in the storage device. If you enable the auto media management feature in the Device resource, the server automatically:

■ Labels the volume

■ Mounts the volume

■ Overwrites volumes it considers unlabeled

■ Recycles volumes eligible for reuse that are loaded into the device.

The auto media management feature is only enabled for stand-alone devices in the Device resource. To enable auto media management for devices in an autochanger, see "Auto Media Management with Autochangers" on page 391.

The Solstice Backup server considers a volume unlabeled with the following conditions:

■ Has no internal label.

■ Is labeled with information other than a recognizable Solstice Backup label.

- Is labeled with a Solstice Backup label, but the density indicated on the internal label differs from the density of the device where the volume is mounted.

Because the auto media management feature can relabel a volume with a different density, it is possible to inadvertently overwrite data that still has value. For this reason, be careful if Solstice Backup volumes are shared between devices with different densities.

If you do not enable the auto media management feature, the Solstice Backup server ignores unlabeled media and does not consider it for backup.

However, if you enable the auto media management feature for a stand-alone device, the following processes occur when a volume becomes full during a backup:

1. A notification is sent that indicates that the server or storage node is waiting for a writable volume. At the same time, the Solstice Backup server waits for the full, verified volume to be unmounted.

2. The device is monitored and the software waits for another volume to be inserted into the device.

3. After a volume is detected, a check is performed to determine whether the volume is labeled. If it is already labeled, the volume is mounted into the device. The Solstice Backup server checks to see whether the newly mounted volume is a candidate to write data to.

   – If so, the write operation continues.

   – If not, the Solstice Backup server continues to wait for a writable volume to continue the backup.

4. If the volume is recyclable and is a member of the required pool, it is recycled the next time a writable volume is needed.

5. If the volume is unlabeled, it is labeled when the *next* writable volume is needed for a save.

If a volume that is not full is unmounted from a stand-alone drive and you enabled auto media management, the Solstice Backup server waits for 60 minutes before it automatically remounts the volume in the drive. This hour delay is designed to give you or an operator time to unload the volume after unmounting.

Caution – The Solstice Backup server considers volumes that were labeled by a different application to be valid relabel candidates if auto media management is enabled. Once the Solstice Backup server relabels the volume, the previously stored data is lost.

Storage Nodes

Storage nodes are host computers with attached remote storage devices. A storage node has the physical connection and ownership of the attached devices, but the Solstice Backup server maintains the client file index and media database. Using Solstice Backup storage node software, you can route client data directly to a storage node's storage devices without the data first going to your Solstice Backup server. For a complete description of storage nodes, see "Storage Nodes" on page 114.

▼ To Configure a Storage Node

You must configure each storage node device. This procedure is similar to the tasks required for adding and configuring devices in "Devices" on page 215. However, the storage node device name must follow this convention:

> **rd=***storage_node_hostname:device_name*

Here is an overview of how to configure a storage node:

1. Install a client (with the required storage node software).

2. Install the driver for each remote device.

3. On the storage node, run jbconfig -s *server* from the command line, where *server* is the host on which the Solstice Backup server resides, to configure the jukebox.

4. Verify that the storage node device name follows this convention:

> **rd**=*storage_node_hostname*:*device_name*

Administrator Privileges

The administrator on the storage node only needs to be on the Solstice Backup server's Administrator attribute to run jbconfig and scanner on the storage node. For normal operation, such as after jbconfig is run on the storage node, the administrator on the storage node can be removed from the server's Administrator list.

> ⚠ **Caution –** It is recommended that you do not attempt to convert an existing Solstice Backup server to a Solstice Backup storage node, as there is a potential for conflict when the resource database, media data base, and client file indexes from the retired server are merged into databases and client file indexes of the controlling Solstice Backup server.

Installing the Software on a New Storage Node

When you install storage node software to create a new storage node, first update your server to the same Solstice Backup release. Then, run the `jbconfig` program on the storage node computer to add a Jukebox resource. After you run the `jbconfig` program on the storage node computer, the Jukebox resource on the server displays values that the `jbconfig` program automatically selects based on the type of autochanger you configure.

Installing an Update Only on the Solstice Backup Server

If you choose to update only the Solstice Backup server with the new storage node software, any existing storage nodes will continue to function as usual, but they will not be able to use the features of the updated `nsrjb` program.

Installing an Update on an Existing Storage Node

If you have existing storage nodes and plan to update both the server and storage node software to this release, first install the updated software packages on the server. Then install the software on the storage node.

▼ To Configure Autochangers and Silos as Storage Node Remote Devices

To configure a remote autochanger or silo device as a storage node remote device:

1. Verify that administrator at the storage node is listed in the Administrator attribute in the controlling Solstice Backup server's Server resource.

   ■ For UNIX storage nodes, use the *root@hostname* format, where `hostname` is the hostname of the storage node.

- For Microsoft Windows storage nodes, use the *administrator@hostname* format, where hostname is the hostname of the storage node.

2. Run the jbconfig program on the storage node to define each device in the autochanger or silo:

```
jbconfig -s servername
```

▼ To Set Timeouts for Storage Node Remote Devices

To set a timeout for a mount request on a storage node remote device and redirect the save to another storage node:

1. From the Media menu, select Devices to open the Devices window.

2. From the View menu, select Details to display the hidden attributes.

3. Set the Save Mount Timeout and Save Lockout attributes to change the timeout of a save mount request on a remote device.

If the mount request is not satisfied within the number of minutes specified by the Save Mount Timeout attribute, the storage node is locked out from receiving saved data for the number of minutes specified by the value of the Save Lockout attribute. The default value for Save Mount Timeout is 30 minutes. The default value for Save Lockout is zero, which means the device in the storage node continues to receive mount requests for the saved data.

Caution – The Save Mount Timeout only applies to the initial volume of a save request.

Storage Management Operations

The following sections detail the various tasks and components involved in the operation of storage volumes and its relationship with pools.

How the Solstice Backup Server Uses Volume Labeling

A volume label is a unique internal code applied by the Solstice Backup server that initializes the volume for the server to use and identifies a storage volume as part of a specific pool. Labeling a volume provides a unique name for tracking and recognizing the media, as well as references to volume labels in the records stored in the media database. The Solstice Backup server uses the media database records to determine which volumes are needed for backing up or recovering data.

When it labels a volume, the Solstice Backup server first verifies that the volume is unlabeled. Then it labels the volume with the name specified in the Volume Name attribute by using one of the following:

- The next sequential label from the label template associated with the chosen pool.
- An override volume name you entered.

If you relabel a recyclable volume from the same pool, the volume label name and sequence number remain the same, but access to the original data on the volume is destroyed. The volume becomes available for new data.

How the Solstice Backup Server Uses Pools

Every volume belongs to a pool. Each pool has a matching label template associated with it. Volumes are labeled according to the rules of these label templates. Label templates provide a method for consistently naming and labeling volumes so there is no need to keep track of how many volumes you have used. You can either take advantage of the preconfigured pools and associated label templates that come with the Solstice Backup product, or create your own pools, label templates, and pool-template associations. Customizing your own label templates gives you more control over your data storage organization.

> **Caution –** Placing a new label on a storable volume, or recycling a volume by relabeling it, makes any existing stored data unavailable for recovery.

How the Solstice Backup Server Selects and Mounts a Volume during a Backup

When a backup takes place, the Solstice Backup server searches for a volume from the appropriate pool to accept the data that needs to be written. The available volumes are those:

- Mounted on stand-alone devices.
- Accessible to the Solstice Backup server through auto media management or via the autochanger.

The results of backup and recovery procedures will vary depending on the availability and condition of the volume:

- If you try to back up files when an appropriate volume is not mounted, the Solstice Backup server requests a writable volume by displaying the following, or similar, message:

```
media waiting (or critical): backup to pool 'Default' waiting for
1 writable backup tape or disk
```

- If you are recovering data, the Solstice Backup server requests the volume name you need to mount in the device by displaying the following, or similar, message:

```
media waiting: recover waiting for 8mm 5GB tape_volume_name
```

- If you need more than one volume to recover your data, the Solstice Backup server displays all the volumes in the order they are needed. During the recovery process, the server requests the volumes it needs, one at a time.

- If more than one volume is mounted on the Solstice Backup server's devices, the server uses the following hierarchy to select a volume on which to write data:

  1. A mounted, appendable volume from the appropriate pool.

  2. A mounted, recyclable volume from the appropriate pool but not currently in use.

  3. An appendable volume that is not currently mounted in the device, but is from the appropriate pool.

  4. A recyclable volume that is not currently mounted in the device, but is from the appropriate pool.

How the Solstice Backup Server Selects a Volume for Relabeling

Backup data is destined for a specific pool. When the data is ready to be written, the Solstice Backup server monitors the active devices to locate a volume from the appropriate pool.

- If only one volume from the pool is mounted and appendable, the data is directed to that volume.

- If two volumes from the same pool are mounted on devices, the Solstice Backup server will consider in the following factors to guide its volume selection:

    - Volume mode

        The Solstice Backup server first writes to a mounted appendable volume from the appropriate pool if one is available. However, if the appendable volume is also recyclable — that is, if all the save sets on the appendable volume have the status recyclable — the volume is treated as recyclable. A mounted recyclable volume from the appropriate pool is the second choice. A mounted recyclable volume that belongs to a different pool will not be recycled. The choice is a new, unlabeled volume.

    - Volume label time (the time when the volume was labeled).

    - Volumes with the oldest label time are selected before volumes that were labeled more recently.

- Number of sessions currently being written to the device.

- If the Solstice Backup server cannot find a mounted volume from the appropriate pool, a mount request is initiated. If auto media management is not enabled or if the server only has stand-alone devices available, mount requests will continue to be generated until a volume is mounted and writing begins.

- If an appropriate volume has been found in the autochanger and more than one empty drive is available, the Solstice Backup server mounts the volume on the lowest-numbered drive. If drives are full, the server will select a loaded idle drive and unload it to make room for the needed volume.

Volume Status

The Solstice Backup server reports on the status of volumes using values such as *Written*, *%Used*, *Location*, and *Mode*. This section defines some of the most common terms used in reports about volumes.

In displays that appear within the Solstice Backup Administrator program, the volume name is the same as the name that appears on the volume label. At the end of the name the following designations may appear:

- (A), which indicates an archive volume
- (R), which indicates a volume that is considered *read-only*

The value of *Written* always indicates the exact number of bytes written to the volume.

The value of *%Used* is an estimate of percentage used based on the total capacity of the volume, based entirely on the specified value of the Media Type of the Device resource. When Used is equal to full, there is no more space on the volume and the end-of-tape marker has been reached or an error with the media has occurred. This is in contrast to a Used value of 100%, which means that the value of *Written* is equal to or exceeds the estimate for this volume.

The volume location refers to an administrator-defined character field that describes a physical location meaningful within the user's environment, such as "2nd shelf, Cabinet 2, Room 42."

For information on changing volume modes, see "To Change the Mode of a Volume" on page 249. For more information about save set status, see "Client and Save Set Combination" on page 153.

TABLE 5-10 lists the Solstice Backup volume modes and their definitions.

TABLE 5-10 Volume Modes

Mode Value	Meaning	Description
appen	appendable	This volume contains empty space. Data that meets the acceptance criteria for the pool to which this volume belongs can be appended.
man	manual recycle	This volume is exempt from automatic recycling. The mode can only be changed manually.
(R)	read-only	The save sets on this volume are considered read-only. The mode can only be changed manually. See the following section for more information on the read-only mode.
recyc	recyclable	The save sets on this volume have exceeded both the browsing and retention policies. The mode can only be changed manually.
full	full	This volume is full; there is no more space for data in the volume, and the save sets have not yet passed the time periods specified by the browse and retention policies.

Read-Only Mode

Changing the mode of a volume to read-only prevents new data from being written to the volume. A read-only volume is *not* a write-protected volume; the save sets on the volume are still subject to their browse and retention policies, and the volume is not protected from being overwritten. When all the save sets on the volume change status to recyclable, the mode of the volume changes to Recyclable, and the volume becomes eligible for automatic recycling.

To change the volume mode, see .

Solstice Backup Volume Operations

If you try to back up files when a volume is not mounted, you will receive one of three messages suggesting that you perform one of the following tasks:

■ Mount a volume

■ Relabel a volume

■ Label a new volume

If you are recovering files, the Solstice Backup server requests the volume name you need to mount. If you need more than one volume to recover the files, the server lists all the volumes in the order they are needed. During the recovery process, the server requests each volume it needs, one at a time. If you back up to an autochanger, the server automatically mounts volumes stored in the autochanger.

▼ To View Volume Information

To view volume information in Solstice Backup Administrator:

1. From the Media menu, select Volumes to open the Volumes window.

Note – You can also click the Volumes button.

---

2. Click on the volume name whose save sets you want to view.

The Volumes window displays the information for the server's volumes, including the volume name, the amount of data written to each volume, the percentage currently used, and other topics.

Check the Volumes window to view the contents of the media database and as a starting point for removing volumes, changing volume modes, and setting locations for your collection of volumes.

TABLE 5-11 lists the volume-related categories displayed in the Volumes window.

TABLE 5-11   Volumes Window

Category	Description
Name	Volume name.
Bar code	Bar code label, if one exists.
Pool	Name of the pool to which the volume belongs.
Written	Amount of data written on the volume.
% Used	Percentage of the volume used compared to its estimated total capacity. When 100% appears in the Used column, the volume has met or exceeded its expected capacity. However, space may still be available because the end of the tape mark has not been reached.
Mode	Status or kind of volume. For more details on the type of volume modes, see "To Change the Mode of a Volume" on page 249.
Location	Location of the volume (entered by the user as a reminder).

▼ To View Save Set Information

To view save set information in Solstice Backup Administrator, select Media>Volumes to open the Volumes window.

Note – You can also click the Volumes button.

The Volumes window lists the save sets stored on a volume. This window displays information, such as the client name, the save set name, the size, and many other categories. You can also change the status of a save set from here.

TABLE 5-11 lists the save set-related categories displayed in the Volumes window.

TABLE 5-12 Volumes Window with Save Set Information

Category	Description
Client	Name of the Solstice Backup client computer that created the save set.
Date	Date when the save set was created.
Level	Level of backup that generated the save set. This refers only to scheduled backups. For manual backups, the level is blank.
Status	Type of save set. For a listing of save set values and descriptions, see "Save Set Status Values" on page 154.
Save Set	Pathname of the filesystem containing the save set. This column also includes clone information. If the save set has a clone, the pathname is marked *has clones* and the cloned save set is marked *clone save set*.

▼ To Mount Volumes

To mount a volume:

1. Insert a volume in the Solstice Backup server storage device or ensure that volumes are already present for the server to access.

2. Select the appropriate device in the Devices list of the Solstice Backup Administrator program.

3. From the Media menu, select Mount to mount the volume.

Note – You can also click the Mount button.

After you label and mount a volume, the volume name is displayed in the Devices list of the Solstice Backup Administration program beside the pathname of the device.

To perform an unattended backup using a stand-alone device, you must mount labeled volumes in the device before leaving it unattended.

⚠ **Caution –** You can only use nonrewinding devices with the Solstice Backup server. If you use a rewinding device, the read/write head is repositioned at the beginning of the volume, and the previously backed-up data is overwritten. The pathnames for these devices must follow the Berkeley Storage Device (BSD) semantic rules, for example, /dev/rmt/0mbn. The "b" in the pathname satisfies the BSD semantics requirement. If you use a file device, you must enter it as a directory path (the same as for other device types) rather than as just a filename.

▼ To Label and Mount Volumes in One Operation

If you have more than one storage device connected to your Solstice Backup server, you must first select the device you want to use for labeling from the Device list in the Solstice Backup Administrator program. Relabeling a volume makes it impossible for the Solstice Backup server to recover original data from that volume.

To label and mount a volume in a single operation:

1. Insert an unlabeled or recyclable volume in the Solstice Backup server storage device, or ensure that these types of volumes are already present for the Solstice Backup server to access.

2. Select the appropriate device in the Devices list of the Solstice Backup Administrator program.

3. From the Media menu, select Label to open the Label window.

Note – You can also click the Label button.

4. Select a pool in the Pool attribute.

   This action also generates a name for the Volume Name attribute. The following rules apply:

   ■ If the volume is unlabeled, the Solstice Backup server assigns the next sequential label from the label template associated with the chosen pool. If you are relabeling a recyclable volume from the same pool, the volume label name and sequence number remain the same. However, access to the original data on the volume is destroyed, and the volume becomes available.

   ■ If you do not select a pool for the volume that you are about to label, the Solstice Backup server automatically applies the label template associated with the Default pool.

   ■ To create individual label names not associated with a template, edit the Volume Name attribute in the Label resource and enter a unique label name.

5. Select the Manual Recycle attribute if you want to manually recycle the volume.

   If you enable the Manual Recycle attribute when you label a volume, the volume cannot automatically be marked as *recyclable* according to the retention policy. When a volume is marked *manual recycle*, the Solstice Backup server disregards the assigned browse and retention policies. Therefore, only an administrator can mark the volume recyclable.

6. Select the Mount After Labeling attribute and click OK. The Solstice Backup server automatically labels the volume, and then mounts the volume into the device.

After a volume is labeled and mounted in a device, the volume is available to receive data. Because the Solstice Backup label is internal and machine-readable, it is a good idea to put an adhesive label on each volume that matches the internal volume label.

For information on using barcode labels, see "Labeling Volumes with Bar Codes" on page 396.

▼ To Label without Mounting

You may prelabel volumes without mounting them. To label a volume without mounting, follow the same procedures as for labeling and mounting in one operation, but deselect the Mount After Labeling attribute in the Label window.

▼ To Unmount Volumes

To unmount a volume:

1. In the Devices list of the Solstice Backup Administrator program, select the appropriate device.

2. Select Media>Unmount to unmount the volume.

Note – You can also click the Unmount button.

Once the volume is unmounted from the device, the Devices list in the Solstice Backup Administrator program displays that information.

▼ To Find a Volume Name

If the adhesive label on the volume is missing or illegible, you can determine its name by using one of these methods:

- Mount the volume in the storage device. The volume name appears next to the device name in Solstice Backup Administrator program. For more information, see .

- Mount the volume in the storage device and enter the nsrmm command at the prompt to determine the volume's name:

```
# nsrmm -pv -f devicename
```

Note – When you use the nsrmm command with the -f *devicename* attribute and your device name contains backslashes, you must use the backslash (escape) character with the backslash. Therefore, a device with the name rd=sn1:\\.\ Tape0 would be written at the command line as: rd=sn1:\\\\.\\Tape0.

- Load the volume in the storage device and select Media>Label. When the Label window appears, the name of the label appears in the Volume Name attribute. Click Cancel to close the dialog without relabeling the volume.

▼ To Change the Mode of a Volume

The mode of a volume determines whether the Solstice Backup server can write to it during a backup. A volume is automatically recyclable if all the save sets on the volume are recyclable. Recyclable save sets have passed the time period designated by the Retention policy for the client.

You can manually change the mode of a full or appendable volume to recyclable if you believe you do not need the data on the volume. The Solstice Backup server then overwrites the data on the volume the next time it needs a volume for a scheduled or manual backup.

To change the mode of a volume:

1. In the Devices list of the Solstice Backup Administrator program, select the appropriate device.

2. From the Media menu, select Volumes to open the Volumes window.

Note – You can also click the Volumes button.

3. From the Volume menu, select Change Mode and one of the following values:

– Appendable

– Read Only

– Recyclable

For a listing of the volume modes, see TABLE 5-10 on page 243.

4. Select a Mode button and click OK.

The changed volume mode is now displayed in the Devices list of the Solstice Backup Administrator program.

Marking a Volume as Full for Off-Site Storage

When removing a volume from your autochanger to store off-site, you should mark the volume as "full" so that the Solstice Backup software will not continue to ask for the volume. To do this, enter the following at the command prompt:

```
# nsrmm -V volid -o full
```

Where *volid* is the volume identifier of the volume. The mminfo command can be used to determine the volume identifier. See the man pages for nsrmm and mminfo or the *Legato Command Reference Guide* for more information about these commands.

Removing Bad Volumes from the Online Indexes

Remove a volume entry from the media database only if the volume has become physically damaged or unusable. When you select a volume for removal, the Solstice Backup server checks to see if the volume has a clone. If a clone is available, the volume entry is not removed from the media database. The Solstice Backup server accesses the clone volume instead of the original volume during recoveries or other actions requiring the volume. If a clone of the volume is not available, then the volume entry will be removed from the media database, and data located on that volume will no longer be browsable.

Removing volume-based entries in the online indexes has the added advantage of reducing the size of the online indexes. However, be cautious when deleting index entries, as it can limit or eliminate your ability to recover data from affected volumes.

The steps to remove volume-based entries are the same whether you are removing a bad volume, reducing the size of your indexes, or both. For instructions on how to perform these tasks, see "Removing Volume-Based Entries from the Online Indexes" on page 129.

Save Set Staging

Save set staging is a process of transferring data from one storage medium to another. Use save set staging to move save sets that have been backed up, archived, or migrated. You can create, edit, and delete staging policies as you can for other Solstice Backup resources. When you enable a staging policy, the Solstice Backup server creates a clone of the save set you specify on a clone volume of the medium you specify. If the save set was stored on a filesystem device, the save set is deleted from the filesystem to free the space. The Solstice Backup server tracks the location of the save set in the media database. The retention policy for the save set does not change when the data is staged.

Staging Data from a Filesystem Device

Staging is particularly well suited for data that has been backed up on filesystem devices. Staging allows the occupied disk space on filesystem devices to be reclaimed so that the disk space can be used for other purposes. Use staging to move the data to more permanent storage, such as an optical or tape volume, or even another, lower-priority filesystem device. Staging also allows data to be moved off the file type device outside the backup period, ensuring that sufficient disk space is available for the next backup session. For more information on backing data to filesystem devices, see "Filesystem Devices" on page 232.

▼ To Create a Staging Policy

To create a staging policy from the Solstice Backup Administrator:

1. From the Customize menu, select Staging to open the Stage window.

2. Click Create.

3. In Name attribute, enter the name of the staging policy.

4. In the Enabled attribute, select Yes to have the staging policy start functioning right away. You have the option of enabling or disabling staging at any time.

5. In the Max Storage Period attribute, enter or select a number. This value indicates the maximum number of days for a save set to be in a volume before it is staged to a different storage medium.

6. In the High-Water Mark (%) attribute, enter or select a number.

   This value is the point at which save sets should be staged, measured as the percentage of available space used on the file system partition that the file device is on. Staging continues until the lower mark is reached (see Step 7).

7. In the Low-Water Mark (%) attribute, enter or select a number.

   This is the point at which staging process should stop, measured as the percentage of available space used on the filesystem partition that the file device is on.

8. From the Save Set Selection attribute, select a save set. This attribute allows you to determine which save sets are selected for staging.

9. From the Destination Pool attribute, select the destination pool for the staged data.

10. Select the filesystem devices that will perform the staging task.

    You must configure the filesystem device before you configure the stage policy in order for the device to appear as a choice. You can assign multiple devices to the staging policy, but a given device should not be controlled by more than one staging policy.

11. In the Start Now attribute, select one of the choices:

    - Recover space
    - Check file system
    - Stage all save sets

Note – The choice you make takes effect immediately. If you do not want any action to start now, leave this attribute alone.

---

12. Click Apply.

▼ To Edit Staging Policies

Note – You cannot change the name of a policy.

To edit a staging policy:

1. In Solstice Backup Administrator, select Customize>Staging to open the Stage window.

2. In the Stage attribute, select the staging policy you want to edit.

3. Edit the attribute values of the staging policy and click Apply.

▼ To Delete Staging Policies

To delete a staging policy:

1. In Solstice Backup Administrator, select Customize>Staging to open the Stage window.

2. In the Stage attribute, select the staging policy you want to delete.

3. Click Delete.

4. When prompted, click OK to confirm the deletion.

5. Click Apply.

Backup Scheduling Tools

This chapter provides information about Solstice Backup software scheduling features, including automated group backups and customizable backup schedules.

This chapter includes the following sections:

Overview of Solstice Backup Scheduling

Scheduled backups of client data is the result of two Solstice Backup resources working in tandem:

- Group
- Schedule

Together, the Group and Schedule resources provide maximum flexibility in scheduling large numbers of differing client backups each day, with the least amount of disruption and inconvenience to your enterprise. The rest of this chapter describes these resources in greater detail.

Group Resource

The Group resource (backup groups) is a resource that specifies a specific time of the day when a backup occurs. Typically backup times are scheduled after regular work hours, such as late evenings and on the weekends. These times are preferred so that the heavy network traffic load associated with backups does not interfere with regular work activities. Any client assigned to a specific group will have its specified save sets backed up during that time.

Schedule Resource

The Schedule resource provides the ability to specify the day of the week and date during the month when the backup occurs. The type and nature of the backup is also specified in this resource. Because scheduled backups are configured in terms of weeks or months (that is, the schedule applies on a monthly or weekly basis), the type of backup scheduled for a specific day of the week is in turn duplicated for:

- The same day in each subsequent week.
- The particular day of a particular week once each month.

Backup Groups

Solstice Backup backup groups allow you to designate a starting time for a client's scheduled backup. Using groups, you can schedule the backups to take place in the middle of the night, or some other time when network traffic is low. You can balance the backup loads by taking the client's backup schedule into account when you decide which clients to include in a specific group. You can also stagger groups at different times of the day to avoid an excessive strain on the network or the Solstice Backup server. For more information about creating schedules that stagger the days that different clients perform full backups, see .

Each client backing up to a Solstice Backup server must be assigned to a backup group, either a preconfigured group or a group you create yourself. In many cases, the Solstice Backup default group is sufficient for the clients, and creating additional groups is not necessary.

When you create a group, the group appears as a choice in the Client resource. The Client resource is where you associate a client with a group. The group name also appears as a selectable option in the pool configuration window, which is where you associate a group with a pool. Remember, you can assign a client's save sets to more than one group.

How the Solstice Backup Server Uses Groups

The client save sets in each backup group begin their automatic scheduled backups according to the start time of the group. You can balance the backup loads by taking the client's backup schedule into account when you decide which clients to include in a specific group.

Any time a Solstice Backup client is added to a Solstice Backup server, either by creating a new Client resource or by adding an existing client, that client is assigned to a backup group. The group determines the time when the assigned client, and any other clients also assigned to the same group, have their data backed up to the server.

Example: Using Groups to Back Up Multiple Clients

FIGURE 6-1 on page 258 illustrates how the Solstice Backup server uses backup groups to back up multiple client save sets. In the example shown, three client computers—mars, jupiter, and saturn—are part of the group named Weekly Full, which starts its automatic scheduled backup at midnight. Client mars runs a full backup of all its save sets every Monday and incremental backups of its save sets on the other days; client jupiter runs a full backup of all its save sets on Tuesday and incremental backups on the other days; and client saturn runs a full backup of all its save sets on Wednesday and incremental backups on the other days of the week. Because each client runs its full backup on a different day of the week, the server is not overloaded.

The second group, *Accounting*, illustrates how you can group clients by department. Group Accounting contains client computers mercury and venus and starts its backups at 7:00 p.m., when the computers in the Accounting Department are available for backup. Although the two client computers run full backups on the same day, computer venus is scheduled to back up only the /usr/home save set; all the save sets on computer mercury are backed up. By estimating how long a backup takes, you can determine what start time to set for the next group.

The save sets from each group are written to appropriate volumes mounted on storage devices. The Solstice Backup server uses pools to organize, track, and store save sets; it uses groups to determine what time clients start their scheduled backups.

Group:	**Weekly Full**		**Group:**	**Accounting**
Daily Start Time:	**24:00 (Midnight)**		**Daily Start Time:**	**19:00 (7:00 p.m.)**

Client:	mars		**Client:**	mercury
Save Set:	All		**Save Set:**	All
Schedule:	Full on Monday		**Schedule:**	Full on 1st of month
	Incremental, all other days			Level 5 on 15th of month
Client:	jupiter			Incremental on all other days
Save Set:	All		**Client:**	venus
Schedule:	Full on Tuesday		**Save Set:**	/usr/home
	Incremental on all other days		**Schedule:**	Full on 1st and 15th of month
Client:	saturn			Incremental on all other days
Save Set:	All			
Schedule:	Full on Wednesday			
	Incremental on all other days			

The Solstice Backup Server writes data to an appropriate storage volume.

(Solstice Backup Scheduled Backup)

Labeled Media Volumes
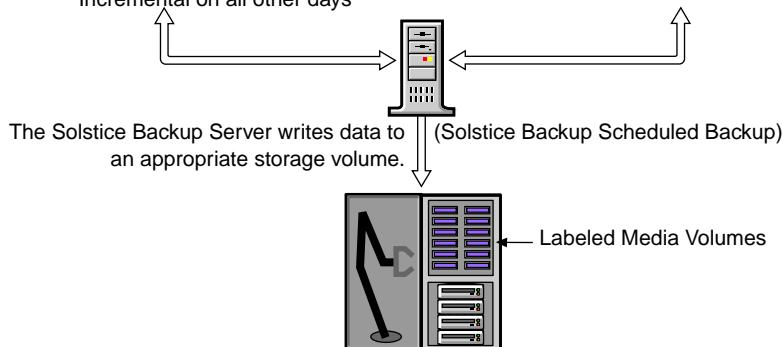
FIGURE 6-1   How Solstice Backup Groups Are Used to Back Up Multiple Clients

Preconfigured Groups

The Solstice Backup product ships with a single preconfigured group named *Default*. To ensure that all data is backed up, the Solstice Backup server automatically adds all clients to the Default group. However, you must enable the Default group for the Solstice Backup server to back it up. Depending on your needs, you can keep a client in the Default group, or you can put the client in one or more customized groups.

You can make changes to any Default group attribute, but you cannot delete the group. You can, however, create or delete as many customized groups as you need.

Key Group Attributes

Depending on your needs, you can keep a client in the Default group, or you can put the client in a customized group.

TABLE 6-1 displays some of the key attributes for the Group resource.

TABLE 6-1   Group Attributes

Attribute	Action
Start Time	The Default group is preconfigured to start its daily backup at 3:33 a.m. This time can be changed.
Interval	This attribute dictates how often a group will start a scheduled backup. The default value is 24:00 (once a day), but you can change this value to start backups more often. For instructions for modifying this attribute, see "To Set the Group Backup Time Interval" on page 261.
Autostart	You must enable the Autostart attribute for the Default group, and any other group you create, before the Solstice Backup server can run a scheduled backup.
Client Retries	If the Solstice Backup server is unable to make a connection with a client, the client will retry connecting up to the value specified in the Client Retries attribute before the backup is considered a failure. The first retry will not occur until after an attempt has been made to at least contact each client in the group.
Inactivity Timeout	This attribute specifies the number of minutes that the Solstice Backup server waits for evidence of backup activity on the client. If the server has not received status information for longer than the time specified, the server abandons the backup operation for the save set.

Abandoned Group Backups

The backup of an abandoned save set might be completed, but the automated report from the savegrp program will not show that the backup is completed. For example, if the client is being backed up over a network file system (NFS) connection and the NFS server crashes and reboots, the Solstice Backup backup hangs until it times out. The Solstice Backup server marks the save set "abandoned," and when the NFS server comes back up, the backup continues and is completed.

▼ To Create a Group

To create a group:

1. In the Solstice Backup Administrator, select Customize>Groups to open the Groups window.

2. Click Create.

3. In the Name attribute, enter the name of the group.

4. For the Autostart attribute, select Enabled.

5. In the Start Time attribute, enter a new time, unless you want to keep the default time of 3:33 a.m.

Note – When you select a start time for your different groups, make sure you schedule them far enough apart so that one group has completed backing up before the next group starts.

---

6. In the Client Retries attribute, change the number of retries (if necessary).

   This value specifies the number of times the Solstice Backup server will attempt to back up a failed client.

7. In the Clones attribute, select Yes to automatically clone the group's backed-up data. The Cloning feature must be enabled for this attribute to work.

Note – If you select No, do not complete the subsequent clone-related tasks in the subsequent steps and skip to step 10.

---

8. In the Clone Pool attribute, select the appropriate clone pool for the cloned data. The Cloning feature must be enabled for this attribute to work.

9. If your environment supports it, select the appropriate migration clone pool in the Migration Clone Pool list. The Cloning feature must be enabled for this attribute to work.

10. In the Printer attribute, enter the name of the printer on which `savegrp` will print the bootstrap save set information.

11. Complete other attributes as necessary, and click Apply.

When you finish creating a backup group, the group appears as a choice in the Client resource where you assign each client individually to a group.

▼ To Set the Group Backup Time Interval

The Solstice Backup server allows you to run an individual scheduled backup group more than once within a 24-hour period. The Interval attribute value of the Group resource determines the frequency (in hours) that an individual group will start a backup.

Note – The Interval attribute is a hidden attribute, meaning that you must display hidden attributes by selecting Details from the View menu before you can access and change the Interval value.

The default value is 24 hours (24:00), which results in one group backup run per day. If you set the Interval attribute value at 12 hours, then the same group will back up twice a day. For example, a group with the default start time of 3:33 a.m. and an interval of 12:00 would back up twice a day, first at 3:33 a.m., and then again twelve hours later at 3:33 p.m.

To set group backup time intervals:

1. In the Solstice Backup Administrator, select Customize>Groups to open the Groups window.

2. From the View menu, select Details to display hidden attributes.

3. Enter the Interval attribute value in the hh:mm format.

4. Click Apply.

   For best results, use time interval values that make it easy to determine the group backup time, such as 24, 12, or 6 hours.

Note – Increasing the group backup time interval (for example, changing the interval from once every 24 hours to once every 12 hours) can add strain to your network, the Solstice Backup server, and associated devices.

▼ To Force an Incremental Backup

For groups that have more than one backup within a 24-hour period, you can use the Force Incremental attribute in the Group resource to instruct the Solstice Backup server to perform only one full or level backup per 24-hour period. If the Force Incremental attribute is set to Yes, the first backup is performed at the configured level. All subsequent backups during the next 24 hours after the start time of the first backup will be incremental, meaning that only changed files will be backed up

regardless of the configured level. This may be advantageous when your environment requires frequent backups, but you want to minimize the impact on systems being backed up.

If the Force Incremental attribute is set to No, multiple full or level backups are allowed during the 24 hours after the start time of the first backup.

To force incremental backups on groups:

1. In the Solstice Backup Administrator, select Customize>Groups to open the Groups window.

2. Select View>Details to display hidden attributes.

3. In the Force Incremental attribute, select Yes.

Note – To disable the force incremental function, select No for this attribute.

---

4. Click Apply.

For more information about configuration group backups that occur more often than once every 24 hours, See "To Set the Group Backup Time Interval" on page 261.

▼ To Edit a Group

You cannot change the name of an existing backup group.

To edit a group:

1. In the Solstice Backup Administrator, select Customize>Groups to open the Groups window.

2. Select the group you want to edit.

3. Edit the attributes of the group and click Apply.

▼ To Delete a Group

To delete a group:

1. In the Solstice Backup Administrator, select Customize>Groups to open the Groups window.

2. Select the group you want to delete.

3. Click Delete.

4. When prompted, click OK to confirm the deletion.

5. Click Apply.

Group Backup Management

Beside providing you with the ability to create and configure backup groups, the Solstice Backup product also provides a variety of tools and features to manage and monitor the operations and progress of group backups.

This section describes how to accomplish the following tasks from the Solstice Backup Administrator program:

- "To Start a Group Immediately" on page 263
- "To Stop a Group" on page 264
- "To Restart a Group" on page 264
- "To View Group Control Details" on page 265
- "To Preview a Group Backup" on page 266

▼ To Start a Group Immediately

When you start a scheduled backup group manually (on demand), each client in the group is backed up. The level of the backup is based on the level for day the manual backup is performed, as defined in the Schedule resource identified in each client's Schedule attribute.

To start a group immediately:

1. In the Solstice Backup Administrator, select Server>Group Control to open the Group Control window.

Note – You can also click the Group Control button.

---

2. In the Groups attribute, select the name of the group you want to start immediately.

3. Click Start to initiate the group backup.

When you start the group in this manner, the Solstice Backup server overrides the groups' scheduled start time and immediately backs up the clients in the group.

▼ To Stop a Group

To stop a group immediately:

1. In the Solstice Backup Administrator, select Server>Group Control to open the Group Control window.

Note – You can also click the Group Control button.

---

2. In the Groups attribute, select the name of the group you want to stop.

3. Click Stop to stop the group backup.

After you initiate the stopping of the group backup, the group stops immediately, any currently running save sessions are aborted, and the Solstice Backup Administrator program displays Not Finished in the Status attribute in the Group Control window.

Caution – If you manually stop a scheduled backup, the Manual Restart option is set automatically, overriding the Autorestart setting. To ensure that the scheduled backup restarts automatically, you should disable the Manual Restart option in the Solstice Backup Administration program.

The Manual Restart option is enabled each time you stop a scheduled backup manually. Consequently, you must disable the Manual Restart option each time you have stopped a scheduled back up manually.

▼ To Restart a Group

To restart a group:

1. In the Solstice Backup Administrator, select Server>Group Control to open the Group Control window.

Note – You can also click the Group Control button.

---

2. In the Groups attribute, select the name of the group you want to restart.

3. Click Restart to reintiate the group backup.

After you initiate the restart of a group backup, the Solstice Backup server resumes the scheduled backup for the group and displays Running in the Status attribute in the Group Control window.

▼ To View Group Control Details

The Group Control Details window displays save sets that:

- Are in the process of backing up.
- Have completed their backups.

Open this window to determine which client's save sets were backed up successfully and which save sets may have failed.

To view group control details:

1. In the Solstice Backup Administrator, select Server>Group Control to open the Group Control window.

Note – You can also click the Group Control button.

---

2. In the Groups attribute, select the group whose details you want to view.

3. Click Details to open the Group Control Details dialog box.

The Group Control Details dialog box displays the status of client save sets in the backup process in one of three message fields:

- Pending Save Sets – displays the client save sets that have not yet been backed up.
- Completed Save Sets – displays the client save sets that the Solstice Backup server has successfully backed up.
- Failed Save Sets – displays the client save sets that the Solstice Backup server did not back up (typically because of a computer or network crash).

  If the server was unable to back up save sets for a client in the group, the backup may have failed for one of the following reasons:

  – The Solstice Backup server failed.
  – The Solstice Backup client failed.
  – Licensing problems.
  – The network connection failed.

▼ To Preview a Group Backup

You can simulate a backup for a specific group by utilizing the Group Control Preview feature. This feature helps you identify potential problems before the Solstice Backup server runs an upcoming group backup. The simulated results are then displayed in the Group Control Details window.

1. In the Solstice Backup Administrator, select Server>Group Control to open the Group Control window.

Note – You can also click the Group Control button.

---

2. In the Groups attribute, select the group whose backup you want to preview.

3. Click Preview to open the Group Control Details window.

Savegroup Completion Report

When the backup is completed, the Solstice Backup server generates a Savegroup Completion Report about the success of the scheduled backup. The server sends the root user an automatic notification and displays the same information in the Solstice Backup Administrator program.

Bootstrap Generation and Printout

When the backup group includes the Solstice Backup server, or if the server is not in an active group, the server generates a special save set called the *bootstrap*, which includes the media database and configuration files. In both of these cases, a bootstrap printout is generated whether the scheduled backup is initiated automatically or manually. The bootstrap information is essential for recovery from a disaster. For information on how the bootstrap is used during a disaster recovery operation, refer to the *Legato Networker Disaster Recovery Guide*.

By default, the bootstrap is printed to the Solstice Backup server's default printer. To change the default printer, edit the Printer attribute in the Group resource.

▼ To Save the Bootstrap to a File

To save the bootstrap to a file:

1. In the Solstice Backup Administrator, select Customize>Notifications to open the Notifications window.

2. In the Notifications attribute, select Bootstrap. The Action attribute displays the following value:

```
/usr/bin/lp -s -c -t bootstrap -d printer_name
```

3. To save the bootstrap information to a file, rather than sending it to a printer, change the Action attribute to the following:

```
/bin/cat >> /directory/filename
```

▼ To Send the Bootstrap File by E-mail

To e-mail the bootstrap file to a user ID, change the Action attribute to the following:

```
/usr/ucb/mail -s nwserver_bootstrap user@hostname
```

System Console Log

The UNIX system log displays messages passed from the Solstice Backup server. When the Solstice Backup server is installed, it adds lines to the configuration log file (`syslog.conf`). These tell the system log facility what types of notices to direct to which file or user. For example:

```
daemon.notice /dev/console
daemon.notice /nsr/logs/messages
daemon.notice operator
local0.notice /nsr/logs/summary
local0.alert root, operator
```

Backing Up Open Files

You can back up most open files using the Solstice Backup product. If an open file changes during a scheduled backup, the Solstice Backup server recognizes that it is changing and backs up the file. The server also checks to see if the size of the file changed. If the file size did change, the server reports the discrepancy. Files that change during backup are displayed in a warning message in the Group Control Details window similar to this:

```
warning: file_name changed during save
```

To make sure the changed file gets backed up, you can do one of the following:

- Restart the backup group.
- Back up the file manually.

Schedules

The Solstice Backup server backs up each client computer across your network according to a backup schedule. A client's backup schedule tells the Solstice Backup server what level of backup operation (for example, full or incremental) to perform on a given day. For instance, on Fridays you might perform a full backup on a client and the rest of the week perform incremental backups.

Schedules can be very simple or very complex, depending on the needs of your environment. All clients can share the same schedule, or each client can have its own unique schedule. The scheduling feature allows you to create customized schedules that you can apply to client save sets. For more information about client configuration, see "Client Configuration" on page 138.

How the Solstice Backup Server Uses Schedules

The Solstice Backup server uses a client's backup schedule to determine what level of backup operation to perform on a given day for specified save sets. The time of day the backup operation begins is determined by the group with which the client save sets are associated.

The type and scope of the backup is determined by the *backup level* you specify. These levels can be set to backup up a client's entire filesystem, or only data that has changed since the last full backup. For more information on how backup levels work, see "Backup Levels" on page 278.

You customize each client's backup schedule to best suit your needs in the Schedules dialog box. For example, some clients may have data you want to back up at the full level every three days, with incremental backups in between. Other clients may have less critical data that only needs a full backup once a month, with incremental backups and/or level 1–9 backups on other days.

Preconfigured Solstice Backup Schedules

Solstice Backup software ships with preconfigured schedules. If these schedules fit your backup requirements, you can use them as is. Otherwise, you can create new ones to accommodate your site-specific needs.

Preconfigured schedules cannot be deleted. Preconfigured schedules that contain "overrides," indicated by an asterisk (*) next to a backup level in the schedule's calendar, cannot be modified. All other preconfigured schedules can be modified. However, a schedule's name cannot be modified; if you want to change the name of a schedule, you must delete the schedule and create a new one.

TABLE 6-2 describes the preconfigured schedules.

TABLE 6-2 Preconfigured Solstice Backup Schedules *(1 of 2)*

Schedule Name	Solstice Backup Backup Operation
Default	Completes a full backup every Sunday, incremental backups on all other days.
Full Every Friday	Completes a full backup every Friday, incremental backups on all other days.
Full on 1st Friday of Month	Completes a full backup on the first Friday of the month, incremental backups on all other days. This schedule cannot be modified.
Full on 1st of Month	Completes a full backup on the first calendar day of the month, incremental backups on all other days.
Quarterly	Completes a full backup on the first day of a quarter; performs a level 5 backup on the first day of the other months in the quarter. Every seven days, a level 7 backup occurs, and incremental backups are performed on all other days. This schedule cannot be modified.
Consolidate Every Friday	Completes a consolidated backup every Friday; completes incremental backups on all other days.

TABLE 6-2   Preconfigured Solstice Backup Schedules *(2 of 2)*

| Schedule Name	Solstice Backup Backup Operation
Consolidate on 1st Friday of Month	Completes a consolidated backup on the 1st Friday of the month; completes incremental backups on all other days. This schedule cannot be modified.
Consolidate on 1st of month	Completes a consolidated backup on the first calendar day of the month; completes incrementals on all other days of the month.
Consolidate Quarterly	Completes a consolidated backup on the first day of each quarter; completes incremental backups on all other days of the quarter. This schedule cannot be modified.

Backup Cycles

The period of time from one full backup to the next full backup is called a *backup cycle*. The following examples demonstrate how to use schedules for different backup cycles and client backup needs. The following items are demonstrated:

- "Example: Weekly Backup Cycle" on page 271
- "Example: Default Schedule with Multiple Clients" on page 272
- "Example: Staggered Weekly Schedules for Multiple Clients" on page 273

Example: Weekly Backup Cycle

FIGURE 6-2 illustrates a weekly backup cycle. In this example, a full backup is performed on a client each Sunday, and incremental backups are performed on the other days of the week.
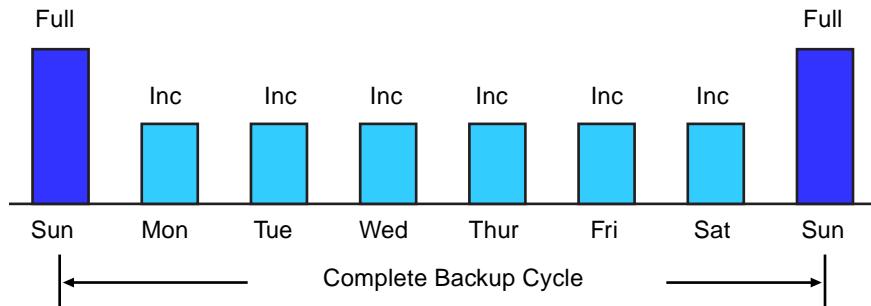
FIGURE 6-2   Weekly Backup Cycle

You can use backup schedules to balance and stagger the load on your Solstice Backup server. Depending on the size of your network, you could apply the same schedule to all clients. For example, if no one works over the weekend and you want to run full backups during this time, you could apply the Default schedule to all of your clients. The Default schedule tells the Solstice Backup server to perform full backups on Sunday, and incremental backups the rest of the week.

Example: Default Schedule with Multiple Clients

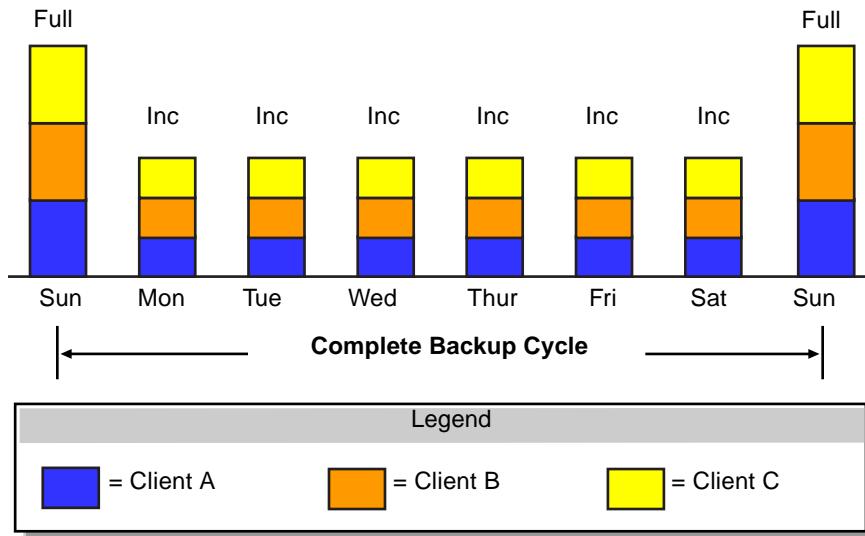illustrates how the Default schedule works for three clients.

FIGURE 6-3   Using the Solstice Backup Default Schedule for Multiple Clients

---

Note – If you have a short backup window period and need to create a full backup, consider using the consolidation backup. A consolidation backup uses the most volume space because it holds both a level 1 and a consolidated full backup. However, data recovery is faster because of the full-level backup. For more details, see "Save Set Consolidation" on page 284.

Since full backups transfer large amounts of data and typically take longer than other backup levels, you may want to stagger them throughout the week. For example, you could apply schedule that performs a full backup for Client A on Thursday, a second schedule that performs a full backup for Client B on Tuesday, and a third schedule that performs a full backup for Client C on Sunday.

Example: Staggered Weekly Schedules for Multiple Clients

FIGURE 6-4 illustrates how to use a staggered backup schedule for multiple clients.
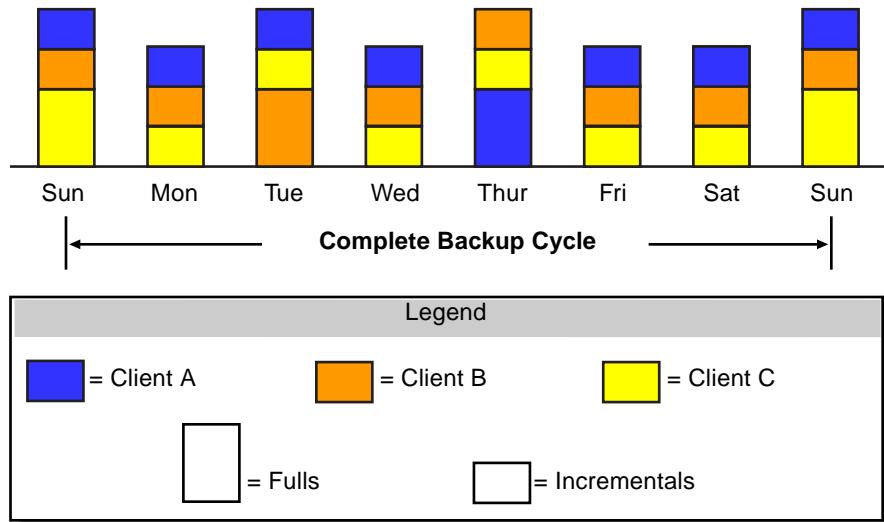

FIGURE 6-4   Staggered Weekly Schedules for Multiple Clients

By balancing and staggering the load and using different start times for different groups of clients, you can increase the efficiency of your Solstice Backup server.

Schedule Planning Considerations

Solstice Backup software makes it easy to set up your backup schedules. Deciding which schedules best fit your environment, however, requires some planning.

When you create backup schedules, consider the following:

- How much data do you have to back up?
- How many backup media volumes do you want to use?
- How much time do you have to complete a backup?
- Does it matter how many volumes are required to recover from a disaster such as a disk crash?

Additionally, you must determine a policy for recovering files. For example, if users expect to be able to recover any version of a lost file that was backed up during a three-month period (that is, the retention policy is equal to three months), you will need to maintain all of the backup volumes for a three-month period. On the other hand, if users only expect to be able to recover data from the last month, you will not need to maintain as many volumes.

The length of time that data is available for recovery by the Solstice Backup server is determined by the browse and retention policies associated with each client. For more information on data life cycle management, see "Browse and Retention Policies" on page 156.

Scheduling Large Client Filesystems

At a moderate backup rate of 400 KB per second, a full backup for a client with 10 GB of data takes about seven hours to complete. Consequently, it might not be convenient to perform a scheduled, full backup for client save sets as large as this because of the amount of time the backup takes.

You can schedule the client's disk volumes for backup at different times by separating them into different backup groups. When you split one client's save sets into multiple backup groups, you back up all the client's files, but not all at once. It is less time-consuming than a full backup of all the local data at one time.

To back up the client's filesystems individually, add and configure the same client several times addressing the different filesystems in the Client resource. For example, configure the first Client resource to back up one filesystem, /usr, with one backup schedule in one group. Then, configure the second Client resource to back up another filesystem, /var, with a second backup schedule in another group.

⚠ **Caution –** When you create separate backup schedules and explicitly list save sets, any files or filesystems not included in an explicit list are omitted from backup. This includes any new disk volumes that are added to the system. This risk of omission does not exist when you enter the special value All in the Save Set attribute. The Solstice Backup software automatically adds the new disk volumes to the backups.

Key Schedule Attributes

To create a customized backup schedule, you must define the following schedule configuration values in the Schedule resource in the Solstice Backup Administrator program.

TABLE 6-3 lists the key attributes contained in the Schedules window.

TABLE 6-3   Key Schedule Attributes

Attribute	Description
Name	The name you give your customized schedule will appear in the Client resource as an attribute you can apply to a client/save set. You should choose a simple, descriptive name; for example, Monday Full. You cannot change the name of an existing schedule. For example, if you want to change the schedule *Full Every Friday* to *Full Every Monday*, you must delete the *Full Every Friday* schedule, and then create a *Full Every Monday* schedule.
Period	The period refers to how often you want a full backup to run. You can set the schedule to apply to either a weekly or a monthly period. • When you select Week and set up a schedule, the backup level Full is applied to that day of the week for all weeks in the calendar year; for example, full backups every Sunday. Week is the default setting. • When you select Month and set up a schedule, the backup level is applied to that day of the month for all months in the calendar year; for example, full backups on the 15th of each month.
Set Level	Use this attribute to set the level of the backup. In addition to full and incremental backups, you can use intermediate backup levels. These allow you to help maintain server efficiency and manage the number of volumes needed for a recovery. You can include one or more of the following levels in a backup schedule: • Full • Incremental • Level [1 – 9] • Consolidated • Skip For more information on how backup levels work, see "Backup Levels" on page 278.
Override	Use this attribute to specify a list of actions and dates overriding an existing backup level for a specific day. For example, you might not want a full backup to run on a holiday. You can override the schedule so the full backup runs on the day before or the day after the holiday.

Note – You might also want to change the Force Incremental attribute value located in the Group resource. The default value for this attribute is Yes. This means an incremental backup will occur if the group is run more than once a day. Set this attribute to No to do more than one full backup per day.

Resource Configuration Order for Schedules

If you want to use your own customized schedule, you must configure the schedule before you can associate it with a client or save set in the Client resource. The start time for your automatic daily scheduled backup is determined by the backup group with which the client save sets are associated. The length of time that the data is available for browsing or recovery is determined by the browse and retention policies you configure for the client's save sets, rather than by the schedule.

▼ To Create a Schedule

To create a schedule:

1. In the Solstice Backup Administrator, select Customize>Schedules to open the Schedules window.

2. Click Create.

3. In the Name attribute, enter the name of the schedule.

4. For the Period attribute, select Week or Month. The value you select changes the way levels set for each day are duplicated:

   - Week: Levels selected for each day of the week are duplicated to the corresponding day of each subsequent week. For example, if you select a full backup for a Friday, every Friday will be configured to have a full backup.

   - Month: Levels selected for each date in the month are duplicated to the corresponding date in the following month. For example, if you select a full backup for the first of the month, every month will have a full backup on the first.

5. Select a backup level for each day of the period (week or month) displayed in the Schedules dialog box by clicking on the day and selecting the backup level from the menu provided. Click Previous Month or Next Month to see the backup calendars for preceding or future months.

   For more information on how backup levels work, see "Backup Levels" on page 278.

6. Click Apply.

You can set up as many Solstice Backup backup schedules as you want, as long as you give each one a unique name.

To assign a backup schedule to a client, you must enter that value in the Schedule attribute of the Client resource configuration.

▼ To Edit Schedules

To edit a schedule:

1. In the Solstice Backup Administrator, select Customize>Schedules to open the Schedules window.

2. Select the schedule you want to edit.

3. Edit the attributes of the schedule and click Apply.

▼ To Delete Schedules

Note – You cannot delete schedules that are currently selected in a client's Schedule attribute.

To delete a schedule:

1. In the Solstice Backup Administrator, select Customize>Schedules to open the Schedules window.

2. Select the schedule you want to delete.

3. Click Delete.

4. When prompted, click OK to confirm the deletion.

5. Click Apply.

▼ To Override a Client's Regular Backup Schedule

You can use a group's Level and Schedule attributes to override a client's regular backup schedule. For example, one evening you might want to run a full backup on all the clients in a group, regardless of the clients' regular backup schedules. The entry you make in the Level attribute overrides the backup level setting for every client in the group.

Alternatively, you might want a group of clients to follow the same backup schedule instead of each client's individual schedule. You could assign a group of clients to follow the default schedule (full every Sunday) regardless of each client's individual schedule. If you leave the group's Level and Schedule attributes blank (the preconfigured setting), the clients follow their individual backup schedules.

To override a regular backup schedule:

1. In the Solstice Backup Administrator, select Customize>Schedules to open the Schedules window.

2. Select the schedule that will contain the override.

3. In the calendar, click on the day to which the override applies.

4. Click on the Override option to open a sub-menu.

5. Select the override level from the menu provided.

6. Click Apply.

Backup Levels

Because it may not be practical or efficient for you to run full backups every day, the Solstice Backup server allows you to specify the level of the backup operation performed during its automatic, scheduled backups. Limiting how often you perform a full backup can help maintain server efficiency, while still ensuring that your data is protected. Different backup levels allow you to trade off the number of volumes and amount of time required to complete a backup with that required to recover from a disk crash.

TABLE 6-4 lists the backup levels supported for filesystem data:

TABLE 6-4   Backup Levels *(1 of 2)*

Backup Level	Function
Full	Backs up all files, regardless of whether or not they have changed.
Level [1 – 9]	Backs up files that have changed since the last backup with a lower or equally numbered backup level. For example: • A level 1 backup backs up all the files that have changed since the last full backup (considered a level zero). • A level 3 backup backs up all the files that have changed since the last level 3 backup, level 2 backup, level 1 backup, or full backup. • A level 9 backs up all the files that have changed since the last level 9, 8, 7, 6, 5, 4, 3, 2, 1, or full (0) backup.
Incremental	Backs up files that have changed since the last backup, regardless of level.

TABLE 6-4 Backup Levels *(2 of 2)*

Backup Level	Function
Consolidated	Backs up all data that has changed since last full backup and subsequently merges these changes with the last full backup. For more information, see "Save Set Consolidation" on page 284.
Skip	Skips the scheduled backup. For example, you may want to skip a backup on a holiday if you know that no one will be available to change or add more media volumes.

Caution – A level Skip backup does not back up any data, but the Solstice Backup server still contacts the client for the scheduled backup at the Start Time indicated for the backup group. The server's savegrp program generates a Savegroup Completion Report that shows that a level Skip backup was performed and no data was backed up. Any problem that could cause the Solstice Backup server to hang waiting for the client to respond has the same effect during a level Skip backup as for other level backups, even though no data is actually backed up.

How the Solstice Backup Server Backup Levels Work

Backup levels work in conjunction with a client's backup schedule. The way you define the backup levels directly affects how long a recovery from a disk crash takes and how many backup volumes you need.

Planning your level backups help you maintain control over the number of volumes you use. A carefully planned backup strategy enables you to recover everything to disk using a minimum number of volumes. The fewer volumes you need to recover from a disk crash, the less time you must spend restoring the disk.

You can also control the size and time it takes to back up your data by using directives, which compress and eliminate unnecessary data from your backups. For example, you can use a directive that tells the Solstice Backup server to skip certain files or filesystems when performing a backup. For more information on directives, see "Directives" on page 178.

The following three examples illustrate the concept of how Solstice Backup backup levels work, with the requirements for recovery in the event of data loss.

Example: Backup Levels Part 1

As shown in FIGURE 6-5 on page 280, a full backup runs on October 1. On October 2, an incremental backup saves everything that changed since the full backup. On October 3, another incremental backup backs up everything that changed since October 2. Then, on October 4, a level 7 backup backs up everything that changed since the full backup on October 1.

To fully recover from a disk crash on October 4, you need only two volumes: the full volume from October 1 and the new level 7 volume. You no longer need the data on the volumes from October 2 and 3, because the level 7 volume includes that information.

Also, incremental backups on October 5, 6, and 7 back up everything that has changed since the level 7 backup on October 4.
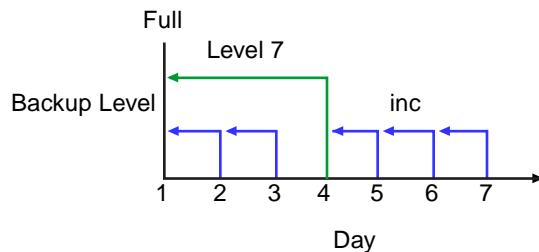
FIGURE 6-5   Backups for October 1 through October 7

Example: Backup Levels Part 2

FIGURE 6-6 continues the example illustrated in FIGURE 6-5 by showing a level 5 backup on October 8, which backs up everything that changed since the full backup on October 1. To fully recover from a disk crash on October 8, you need only two volumes: the full volume from October 1 and the new level 5 volume. You no longer need the data on the volume from the level 7 backup on October 4 (or the subsequent incremental backups performed on October 5, 6, and 7) because the new level 5 volume includes that information.
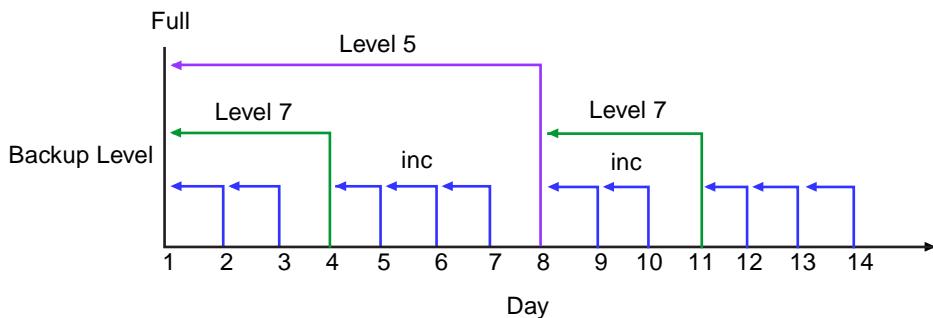
FIGURE 6-6   Backups for October 1 through October 14

Also, on October 11 a level 7 backup backs up all of the data that changed since the last lower-numbered backup (in this case, the level 5 backup on October 8). To recover from a disk crash on October 11, you need three volumes: the full volume from October 1, the level 5 volume from October 8, and the new level 7 volume.

Example: Backup Levels Part 3

Continuing the example, in FIGURE 6-7 a level 5 backup on October 15 backs up all of the data that changed since the last lower-numbered backup. Because no backup lower than level 5 has been performed since the full backup on October 1, this level 5 backup backs up all of the data that changed since the full backup. Therefore, to recover from a disk crash on October 15, you only need two volumes: the level full volume from October 1 and the new level 5 volume.
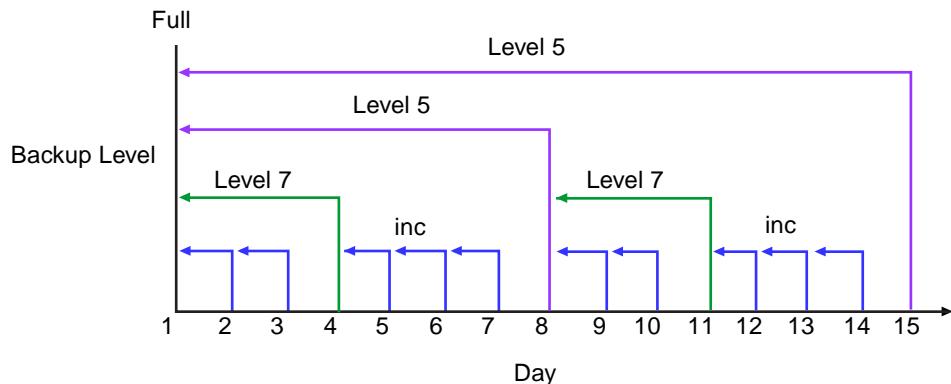
FIGURE 6-7   Backups for October 1 through October 15

How the Solstice Backup Server Uses Backup Levels

A backup schedule defines what level backup the Solstice Backup server should perform on a given day during a backup cycle. You can apply one or more of these backup levels to customize a backup schedule. If you are considering using backup levels in a customized schedule, consider the following to help you decide what best suits your environment:

■ Full backups take more time to complete than incremental backups.

■ If you have only one storage device and the full backup does not fit on a single piece of media, an operator must be available to monitor the backup and change the media.

■ Full backups cause the online indexes to grow more rapidly than incremental or level backups.

■ Level backups serve as checkpoints in your schedules because they collect all the files that have changed over several days, or even weeks, into a single backup session.

■ Consolidated backups provide the same benefits at the same cost as do full backups. The essential difference, however, is that consolidated backups are less taxing on the network and client because only a level 1 backup is performed. The server's performance, however, is slower because the server consolidates the changed data with the original backup.

TABLE 6-5 lists advantages and disadvantages of each backup level.

TABLE 6-5   Advantages and Disadvantages of Backup Levels

Backup Level	Advantages	Disadvantages
Full	• Faster recovery	• Slow backup • High server load • High load on client and network • Uses more volume space
Level	• Faster backup • Low load on server • Uses least volume space	• Slow recovery • Data can spread across multiple volumes
Consolidated	• Faster backup (from the client's perspective only) • Faster recovery • Low load on client and network	• Longest high load on server • Requires at least two volume drives • Uses most volume space

Online Indexes and Backup Levels

The online indexes are backed up in the following manner:

■ The client file index is backed up every time the client is backed up, and is backed up at the same level as the client.

■ The server index is backed up every time the server is backed up, and it is backed up at the same level as the server.

■ The media database is backed up whenever the server is backed up, or after every scheduled backup if the server is not in an active group.

For example:

■ If the Solstice Backup server is backed up at the level full, the backup levels of the server index and media database are also full.

- If the Solstice Backup server's backup is a level 5, the backups of the server index and media database is also a level 5.

  However, when the server's backup level is incremental, the backups of the server index and media database are level 9. For a consolidated backup, the backups of the indexes are full-level backups. The server does not perform a level 1 backup for this data.

Save Set Consolidation

This section describes the *save set consolidation* feature.

How Save Set Consolidation Works

Save set consolidation merges a new level 1 backup with the last full backup of a save set to create a new full backup. Essentially, save set consolidation is a process rather than an attribute; it describes the process of merging a level 1 backup with an existing full level backup. Consolidation is really a full level backup and the resulting save set of the consolidation process is the same as a full level save set. Although a consolidated backup takes place, there is really no such thing as a consolidated save set.

As such, you cannot find *consolidation* as a backup level in most of the various resources where Solstice Backup backup levels appear. For example, you cannot find *consolidation* as a potential backup selection in the Pool resource. The one exception is the Schedule attribute, where you can use the default consolidation schedules provided, or create your own consolidation schedule by overriding a scheduled level with the special level *c*.

Save set consolidation eliminates the need to perform full backups at regular intervals. After scheduling a save set consolidation backup, you only perform one full backup during the first scheduled consolidated backup cycle. Afterward, all subsequent backups of the consolidated save set are incremental backups.

Save set consolidation takes place entirely on the server. The burden of backups is shifted entirely to the server, reducing Client resource use and network traffic. This shift to the server also results in more frequent level saves, resulting in fewer tapes to process in the event of a full recovery.

When working with large files, save set consolidation shortens the backup window by doing incremental saves, thus reducing the number of tape drives required during a recovery.

Note – If there are no existing full backups the first time a consolidation backup begins, the consolidation backup defaults to a level of full. Once a full backup exists, subsequent consolidation backups are performed as a level 1 followed by consolidation with the existing full.

Requirements

To initiate save set consolidation by the server, the client must be running a Legato ClientPak release 5.5 or later. In release 5.5, the administrative interfaces within the Legato ClientPak software were updated to provide for initiation of save set consolidation. Save set consolidation will work for save sets of clients whose Legato ClientPak release is earlier than 5.5, but the consolidation must be initiated by the Solstice Backup server itself.

Save set consolidation requires at least two attached devices. For better performance, you should have three or more devices available for consolidated backups.

Save Set Consolidation Usage

You should use save set consolidation if the following conditions are present:

- A client is at a remote location and data transfer over the network to the server is a performance issue for either the network or the client.
- Either your network bandwidth is small, or large backups over the network are cost-prohibitive.
- You need to back up large filesystems with little incremental data.
- The server has the necessary resources (a minimum of two volume drives and preferably three or more volume drives) and the workload capacity to consolidate full backups locally.

Note – Do not consider save set consolidation as a performance improvement without considering the above-listed conditions. Under the appropriate conditions, save set consolidation can be used to free network bandwidth and client resources. Used inappropriately, save set consolidation can lower performance, since it uses tape-to-tape transfer, which might be slower than the disk-to-tape data transfer used by backups.

Limitations of Save Set Consolidation

There are some limitations to the save set consolidation process. In the following circumstances, save set consolidation will fail:

- Raw disk file partitions are not eligible for consolidation.

- Database systems cannot be consolidated.

- Renamed and deleted directories cannot be consolidated. If incremental data contains renamed or deleted directories, the save set consolidation process will detect it and abort.

- You can only administer the server performing the save set consolidation using the graphical Solstice Backup Administrator program from a Solstice Backup release that includes the save set consolidation feature.

In addition, we recommend you do not use save set consolidation if any of the following conditions are present:

- The client is connected to the server over a fast network or is a local client, and the network data traffic generated by full backups is not a problem. In this instance, save set consolidation will not produce a measurable benefit.

- The filesystems being backed up are either not very large or they contain a large number of small files that are modified often.

- The incremental data usually contains a large amount of data, and the number of files that have changed since the last full backup is large compared to the total number of files in the filesystem.

- It is cost-prohibitive to allocate three (or the minimum of two) tape drives for the exclusive use of the server while it consolidates the full backup.

> **Caution –** Even if a consolidated backup cannot be completed because of a system crash, tape drive errors, or other problem, data integrity is ensured. The consolidated backup will be aborted. The transactional log assists save set consolidation in putting back changes made to the online index. In most instances, the Solstice Backup server only makes online index changes when data has been written to tape.

Invoking Save Set Consolidation

You can schedule save set consolidation through the Solstice Backup Administrator program or invoke it from the command line.

▼ To Schedule Save Set Consolidation

To schedule save set consolidation through the Solstice Backup Administrator program:

1. In the Solstice Backup Administrator, select Customize>Schedules to open the Schedules window.

2. In the calendar, click on the day that you want the save set consolidation to occur.

3. Select Consolidate from the menu provided.

The letter *c* now appears for that day, signifying the newly created consolidated backup. The Solstice Backup server will invoke the save set consolidation backup on this date.

For a consolidated backup, the Solstice Backup server will first perform a level 1 backup; subsequently, it performs a consolidation of the level 1 backup and the most recent full save set.

▼ To Invoke Save Set Consolidation from the Command Line

You can invoke save set consolidation through two different command-line programs:

- savegrp

  The savegrp program must be used with the -lc option to indicate that the backup level is consolidated, or with a schedule that has a level *c* on the schedule. This program automatically performs a level 1 backup, followed by the actual consolidation process.

- nsrssc

  The nsrssc program completes the consolidation process. For the nsrssc program to be successful, a level 1 save set and a level full save set must *already* exist.

Note – When nsrssc is run from the command line, bootstrap information will not be written to the media used for consolidation. As a result, you will still need the original media for disaster recovery.

Both programs also offer other options to maximize the flexibility of save set consolidation. For further information, refer to the `savegrp` and `nsrssc` man pages or the *Legato Command Reference Guide*.

▼ To Direct Data from a Consolidation Backup to a Specific Pool

By default, save sets from a consolidation backup are written to whatever media is mounted for the group most recently backed up.

To direct consolidated save sets to a specific set of media:

1. Configure a Group resource for consolidation backups. See "To Create a Group" on page 260.

2. Configure a Pool resource for consolidation backups. See "To Create a Pool" on page 204.

3. Select the name of the Group resource you created in step 1 as the value for the Groups attribute.

4. Edit the Client resource for each client for which you want to perform consolidated backups, and assign the client to the group that you created for consolidation backups.

Example: Consolidation of Level 1 and Full Backups

FIGURE 6-8 illustrates the concept of how the Solstice Backup server merges a level 1 backup with an existing full level backup to "consolidate" both to create another full level backup.

On October 16, the consolidated level "c" automatically performs a level 1 backup, which backs up all the data that has changed since the last full backup. As part of its consolidation process, the Solstice Backup server merges this level 1 backup with the last full backup (created on October 1) and builds a new full level backup. To recover from a disk crash on October 17, you need one volume, the full volume that was created using the consolidation process on October 16.
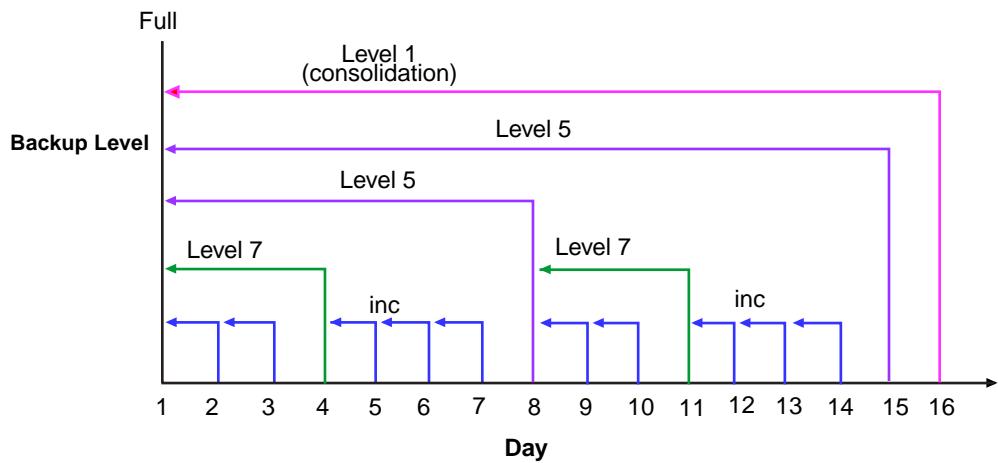
FIGURE 6-8   Consolidated Backup

Data Backup and Recovery

This chapter describes the method for backing up your data manually, rather than using the scheduled group backup method. It also details the various methods for recovering backed-up data, and how to direct recovered data to the location of your choice.

The main sections in this chapter are:

Manual Data Backups

The Solstice Backup server backs up your client data regularly using scheduled backups. Scheduled backups are the preferred backup method because they occur automatically, and also because scheduled backups can be recovered more easily.

Manual backups are intended for quick backups of a few files initiated by a user. They do not generate bootstrap files or back up indexes, so they cannot replace scheduled backups. Remember that you can manually start a scheduled backup group at any time by selecting the Start Group command. For a complete description of scheduled backups and groups, see Chapter 6.

There may be times, however, when you want to back up one or more data items to the Solstice Backup server right away, instead of waiting for the scheduled nightly backup. For example, if you have just finished working on a spreadsheet, you might like to back up the file right away.

▼ To Manually Back Up Data

The `nwbackup` program provides a graphical interface through which you can perform manual backups. To view a display of the `nwbackup` program, see "The nwbackup Program" on page 84.

The `nwbackup` program is used to perform a manual backup initiated from the client rather than from the server's scheduled group backup (`savegrp`) that includes the client.

You can designate the entire filesystem, specific directories, or individual files for manual backup. If you select a directory for backup, all of its nested subdirectories and files are also selected for backup. If you want to exclude certain files from manual backup, you must deselect those items.

You can also search the filesystem for files you would like to back up by selecting View>Search to open the Search window and entering the name or a partial name of the item you are seeking.

To start a manual backup:

1. Start the `nwbackup` program by entering the following command at the prompt:

```
# nwbackup
```

Use the `-s` *server* option with `nwbackup` to specify a particular Solstice Backup server on the network.

The `nwbackup` program displays a representation of the client's filesystem.

2. Select each directory and/or file you intend for manual backup by performing one of the following tasks:

   ▪ Select the directory or file and click the Mark button

     To clear an item, click the Unmark button.

   ▪ Select the directory or file and select Mark from the Mark menu.

     To clear an item, select Unmark from the Mark menu.

   ▪ Select the check box next to each directory or file listed.

     To clear a check box, click the check mark next to the selected item.

When you mark a directory or file for backup, a check mark appears next to that item.

3. Click the Start button to begin the manual backup.

   When the Solstice Backup server has successfully finished the backup, the `nwbackup` program displays the following (or a similar) message:

```
Backup completion time: 2-15-97 3:27p
```

Caution – Sometimes there are bad files or errors on computer disk volumes that the Solstice Backup server cannot detect during a backup. It simply backs up the bad data. It is a good idea to run diagnostic programs regularly on your computers to correct disk volume errors.

Data Recovery

Recovery of backed-up data is performed using the `nwrecover` program. This program is similar to the `nwbackup` program discussed in the previous section. To view a display of the `nwrecover` program, see "The nwrecover Program" on page 85.

When the Solstice Backup server backs up a file to media (for example, magnetic tape), it creates an entry for the file in the client file index. The Solstice Backup server uses this index to map the backed-up files to volumes. If your computer has not been backed up yet, there will be no entries in your client file index, and therefore no backups are displayed in the `nwrecover` program's main window.

Using `nwrecover`, you can browse backed-up files and filesystems from Solstice Backup clients and recover files as they were at a specific point in time when they were backed up. The versions of files and filesystems available in the `nwrecover` are determined by the time period specified in the browse policy. For more information about the browse policy, see "Browse and Retention Policies" on page 156.

If remote access is enabled, users may also be able to recover files from other client computers. For more information on remote access, see "Remote Access Rights for Other Clients" on page 146.

When recovering data onto a filesystem where other versions of the same data already reside, there is the likelihood of a file conflict. You can determine what course of action to take each and every time a conflict occurs during a recovery, or

you can instruct the Solstice Backup server in advance what action to take for all conflicts. You can also simply redirect the recovered data to a different directory where there is no danger of conflicts.

▼ To Recover Data

You can designate the entire filesystem, specific directories, or individual files for recovery. Note that if you select a directory for recovery, all of its nested subdirectories and files are also selected for recovery. If you want to exclude certain files from recovery, you must deselect those items.

To recover data:

1. Start the `nwrecover` program by entering the following command at the prompt:

```
# nwrecover
```

- Use the `-s` *server* option with `nwrecover` to specify a particular Solstice Backup server on the network.
- Use the `-c` *client* option to specify a particular Solstice Backup client on the network. If the `-c` option is not used, the current client is assumed.

The `nwrecover` program displays a representation of the client's filesystem.

2. Mark each directory and/or file you intend to recover by performing one of the following tasks:

- Select the directory or file and click the Mark toolbar button

  To clear an item, click the Unmark button.

- Select the directory or file and select Mark from the Mark menu

  To clear an item, select Unmark from the Mark menu.

- Select the check box next to each directory or file listed.

  To clear check box, click the check mark next to the selected item.

When you mark a directory or file for recovery, a check mark appears next to that item.

Note – If you want to recover data from a scheduled backup other than the most recent one, then you must change the browse time. See "To Recover Data from a Previous Backup Time" on page 295.

---

3. Click the Start button to begin the recovery.

It takes the Solstice Backup server a few moments to recover your file, depending on file size, network traffic, server load, and tape positioning. During this time, messages appear for each recovered file in the nwrecover program, so you can monitor the progress of the recovery.

If there is a naming conflict with the file you are recovering, the Solstice Backup server prompts you to choose one of several workarounds. For more information about naming conflicts, see "Resolving Recovered Data Conflicts" on page 296.

If the recovery is successful, the nwrecover program displays the following (or a similar) message:

```
Successfully received 1 file

Recovery completion time: 5-30-97 2:03p
```

▼ To Recover Data from a Previous Backup Time

Using the nwrecover program, you can browse the entries for each backed-up file listed in the client file index created at any time you specify. Changing the browse time allows you view the entries for your backed-up files for a particular period. You might want to do this if you need to:

■ Look for data that has been deleted from your file system and may not be present in the most recent backup.

■ Retrieve an earlier version of existing data.

Your browse policy determines how long older data is retained and thus, your ability to browse and recover this data. For example, if you would like to recover a lost file that you backed up six months ago, your browse policy must support it. For more information, see "Browse and Retention Policies" on page 156.

To recover data from a previous backup time:

1. In the nwrecover program, select Change>Browse Time to open the Change Browse Time window.

   The Change Browse Time window displays a monthly calendar in which you can change the date and time of the backup being displayed. The Change browse time window helps you find files that have been deleted and are not in the most recent backup, or allows you to recover earlier versions of existing files.

2. Select the month and day that corresponds to the client data you want to recover.

3. In the Browse Time attribute edit the time and date listed to change what is displayed.

4. Click OK.

The Change Browse Time window disappears, the nwrecover program now displays your client's filesystem as it appeared on the day and date you selected. The nwrecover program also displays that date in the Browse Time attribute.

5. Select the files or directories you want to recover as described in "To Recover Data" on page 294.

Note – If there are no entries for files backed up at the time selected in the Browse Time field, the Solstice Backup software displays files that were backed up at the closest preceding time. However, if the time selected in the Browse Time field is earlier than the first entry in the client file index, the error "Notice! Nothing in index at this time" is displayed.

Resolving Recovered Data Conflicts

By default, when the Solstice Backup server recovers data it will attempt to return them to their original location, using their original file and directory names. However, if another file or directory with the same name is present in that same location, a file recovery conflict will occur. You can instruct the Solstice Backup server to respond to these recovery conflicts in several different ways, either on an file-by-file basis, or automatically.

You are prompted to make conflict resolution decisions each time to start a recovery using the nwrecover program. When you start the recovery, the Conflict Resolution window appears and asks if you want to be prompted when a conflict occurs.

- If you select Yes in the Conflict Resolution window, then each time there is a conflict, the Solstice Backup server will prompt you to select a resolution option.
- If you select No in the Conflict Resolution window, then you must enter an automated conflict resolution option that the server will follow each time there is a conflict during the recovery.

TABLE 7-1 lists the conflict resolution options available during a recovery.

TABLE 7-1   Recovery Options

Option	Action
Rename the Recovered File	Renames the recovered files by appending each conflicting name with a .R.
Discard the Recover File	Discards the recovered file in order to prevent the Solstice Backup server from overwriting an existing file.
Overwrite the Existing File	Replaces the file on your filesystem with the recovered version.

You can also choose to relocate your recovered data to a different directory where you know there will be no conflicts, or create a new directory using the Relocate feature. However, even if you relocate your recovered data to a new and/or empty directory, the Solstice Backup server still prompts you to enter your conflict resolution information. For more information on this feature, see "To Relocate Recovered Data" on page 297.

▼ To Relocate Recovered Data

When you recover a file, by default the Solstice Backup server will attempt to copy the file to it original location. Unless you relocate the recovered file, the server tries to place it in the original directory where the file resided when it was backed up. If an existing file on your local disk has the same name as the one being recovered, the server prompts you for instructions on how to recover the file (rename, overwrite, or discard).

You can instruct the Solstice Backup server where to relocate the recovered file in advance of the recovery using the Relocate feature in the `nwrecover` program. Enter the new recovery destination path in the Relocate To attribute of the Recover Options dialog box.

To relocate recovered data:

1. In the `nwrecover` program, select Options>Relocate to open the Relocate window.

2. In the Relocate To attribute, enter the new directory location and click OK.

   You must provide the full pathname of the relocation directory.

3. Perform a data recovery as described in "To Recover Data" on page 294.

The recovered data will automatically be relocated to the directory you entered.

▼ To Search for a Recoverable File

Users can locate files or directories for recovery using the `nwrecover` program's Search feature by entering the name of the item you are seeking and the directory where the search will begin. The `nwrecover` program then displays the first file that matches the search pattern.

To search for a recoverable file:

1. In the `nwrecover` program, select View>Search to open the Search window.

2. Enter the name of the file or directory. Wildcards can be used to expand the search.

3. Enter a path to designate where the search will begin.

▼ To View Details of Recoverable Files

You can display the contents and other details of directories listed in the left pane of the `nwrecover` window in the right pane using the File Details feature.

When the File Details feature is enabled, the `nwrecover` program displays detailed information regarding the files and subdirectories of the selected directory, including the date and time the file was last backed up, the permission of the file, and the date created.

To view details of recoverable files:

1. In the `nwrecover` program, select the item whose details you want to view.

2. Select View>File Details to display this information.

▼ To View the Version Details of Backed-Up Data

You can view the back-up details of a directory or file you select in the `nwrecover` program using the Versions feature. The Versions feature provides a listing of each time the Solstice Backup server backed up that file, including the day, date, time, and size.

To view the version details of backed-up data:

1. In the `nwrecover` program, select the item whose backup version you want to view.

2. Select View>Versions to open the Versions window.

▼ To View Volumes Required for Data Recovery

You can view a list of the volumes needed to recover files and directories you have selected for recovery in the `nwrecover` program using the Volumes feature. The Volumes feature displays a window with the names of all the volumes the Solstice Backup server needs to access for the data recovery.

To view a list of the volumes required for data recovery:

1. In the `nwrecover` program, select the item whose volume information requirements you want to view.

2. Select View>Volumes to open the Recover Volumes Required window.

Disaster Recovery

Recovering data from Solstice Backup servers damaged or disabled as a result of a disaster or other catastrophic event requires more complex steps than are covered in this chapter. Optimum recovery also requires careful planning. For more information, refer to the *Legato NetWorker Disaster Recovery Guide*. For the most recent disaster recovery information, refer to the *Sun Solstice Backup Release Supplement.*

Note – You should keep a printed copy of this documentation in a safe, convenient place away from the server.

Directed Recovery

Directed recovery allows you, from an administering Solstice Backup client at a remote location, to direct a Solstice Backup server to perform both of the following:

■ Retrieve backed-up data from one of its Solstice Backup clients (a source client)

■ Recover that data to a different Solstice Backup client (a destination client) within your enterprise.

The destination client receiving the recovered data can be the source client itself, or any other Solstice Backup client within the enterprise, including the administering Solstice Backup client. Because you do not have to be physically present at either the source client, the destination client, or the Solstice Backup server itself, you have a high degree of flexibility and convenience in recovering and redirecting data. Directed recovery includes all of the Solstice Backup server's standard file handling features.

You use the Solstice Backup command line `recover` program to administer directed recoveries. The `nwrecover` program does not support directed recovery tasks.

Uses of Directed Recovery

Directed recovery is very useful when data recovery is required for a shared server or other computer where users cannot recover the data themselves. If you have a Solstice Backup client installed and the appropriate access rights, the directed recovery feature allows you to perform the recovery from your own computer remotely rather than from the source computer itself.

Directed recovery also permits the recovered data to be directed to other Solstice Backup client computers besides the requesting client computer. This allows you to share files and information with a broad range of Solstice Backup client computers.

Access Requirements for Directed Recoveries

Directed recovery is a Solstice Backup administrator-level function and is typically not accessible for users and other nonadministrators. Performing directed recoveries requires you to have administrative access for both the Solstice Backup server storing the data from the source client, and all destination clients receiving the recovered data. If you anticipate performing directed recoveries on a regular basis, you should establish the required permissions for each Solstice Backup server and client in your enterprise prior to receiving data recovery requests.

Access requirements to perform directed recoveries of UNIX-based Solstice Backup clients are as follows:

- Both the Solstice Backup client administering the directed recovery and the destination client must be release 6.0 or later, because Solstice Backup software releases earlier than 6.0 did not support directed recovery on UNIX platforms. The source client and the Solstice Backup server can use Solstice Backup software earlier than release 6.0, as long as the save sets use the earlier than release 6.0 format.

- Your user name must be included in the Administrator List attribute of the Solstice Backup server that backs up the source client's data; for example, *username@administration_client*.

- The hostname of the Solstice Backup client with which you administer the directed recovery must be listed in the servers file of each Solstice Backup client computer designated as a potential destination client.

- Users that do not have Solstice Backup administrator rights to the Solstice Backup client administering the directed recovery and the destination client must be included in the Remote Access attribute of each client, for example, `user@administering_client_hostname` and `user@destination_client_hostname`. If you already have Solstice Backup administrator rights, you automatically have remote access rights.

The following sections describe how to set the various server and client administration rights required to enable directed recovery.

▼ To Grant Administrative Rights to a Remote Solstice Backup Server's Client

To grant administrative rights to a remote Solstice Backup server:

1. Start the Solstice Backup Administrator program.

2. Select Server>Server Setup to open the Server window.

3. Enter your `user@administering_client_hostname` in the Administrator attribute and click Apply.

4. Enter the `nsr_shutdown` command at the prompt to shut down the Solstice Backup daemons. A list of Solstice Backup daemons that will be shut down appears, along with a prompt that asks whether you want to continue with the `nsr_shutdown` command as follows:

```
# nsr_shutdown

nsr_shutdown will kill the following processes
 25768 ?          S   0:02 /usr/sbin/nsr/nsrexecd
 25770 ?          S   0:01 /usr/sbin/nsr/nsrexecd
 25771 ?          S   0:02 /usr/sbin/nsr/nsrd
 25783 ?          S   0:00 /usr/sbin/nsr/nsrmmdbd
 25784 ?          S   0:00 /usr/sbin/nsr/nsrindexd
 25785 ?          S   0:00 /usr/sbin/nsr/nsrmmd -n 1
Do you want to continue? [Yes]? y
```

5. Open the `servers` file, located in the `/nsr/res` directory, using your favorite text editor.

6. Enter the administering Solstice Backup client computer's name in the `servers` file using the following format style:

```
mercury.jupiter.com
```

7. Save and close the `servers` file.

8. Start the Solstice Backup daemons by entering the following at the command prompt:

```
# /etc/init.d/networker start
```

▼ To Grant a Network Connection to a Destination Client

The administering Solstice Backup client computer must have access to the destination client in order to recover files onto it. You must add the administering Solstice Backup client computer name to the *servers* file of that destination computer.

To grant remote recovery permission to a destination client:

1. On the destination client computer, open the `servers` file, located in the `/nsr/res` directory, using your favorite text editor.

2. Enter the administering Solstice Backup client computer's name in the *servers* file using the following format style:

```
mercury.jupiter.com
```

3. Save and close the *servers* file.

4. Shut down the Solstice Backup daemons by entering the following command at the prompt:

```
# nsr_shutdown
```

A list of Solstice Backup daemons that are shutting down appears. When prompted, enter Y to confirm that you want to shut down the daemons.

5. Start the Solstice Backup daemons by entering the following at the command prompt:

```
# /etc/init.d/networker start
```

▼ To Gain Remote Access to a Solstice Backup Client

The administering client and the destination client must have remote administration access privileges for any source Solstice Backup client involved in a directed recovery. This access allows you to browse the source client file index and retrieve its files. You can add your *hostname* or *user@hostname* to the remote access list of any Solstice Backup client through the Solstice Backup Administrator program. Note that

if you set the default client of a Solstice Backup server to include your hostname in its Remote Access attribute, all subsequently added clients will also list you as having remote access.

To gain remote access to a Solstice Backup client:

1. Log in to the Solstice Backup server which includes the Solstice Backup client to which you want remote access.

2. Start the Solstice Backup Administrator program.

3. Select Clients>Client Setup to open the Clients window.

4. Select the client to which you want remote access in the Clients attribute.

5. Enter the userid and `user@administering_client_hostname` of the administering client in the Remote Access attribute.

6. Enter the Solstice Backup client daemon account in the Remote Access attribute (`root@`*destination_client_hostname*) for every destination client you need or anticipate needing to access this source client resource.

Directed Cross-Platform Recovery

Directed cross-platform recovery allows you to recover backed-up data remotely from a Solstice Backup client (the source client) and direct it to another Solstice Backup client (the destination client) within an enterprise. The destination client receiving the recovered data can be the source client itself, another Solstice Backup client within the enterprise, or the administering client computer.

You can perform directed recoveries of data between different combinations of Windows-based and UNIX-based source and destination clients. You must, however, have established the proper remote access requirements for the clients and the Solstice Backup server that backs up their data. The source and destination clients must also be of the same type (either both Windows-based, or both UNIX-based).

In the Solstice Backup 6.0 and later releases, the save stream format is different from the format in previous releases. Therefore, you cannot use an earlier Solstice Backup release to recover data that was backed up using a Solstice Backup 6.0 or later release.

Performing a Directed Recovery

You perform directed recoveries between a UNIX source and a UNIX destination client using the **recover** command and various options at the command prompt. The **nwrecover** program does not support directed cross-platform recovery. Remember that a directed recovery is an interactive process only. Once you invoke the recovery, you still need to interactively select and recover files you choose from the source client index.

The `recover` Command in Directed Recoveries

The **recover** command has a variety of options you can add to modify the basic browse and recover functionality. For performing directed recoveries, here are the following options:

TABLE 7-2   Recover Options

Option	Meaning
–d *destination*	The name of the directory where you wish to relocate the recovered files.
–s *server*	The name of the source client's backup server.
–c *client*	The name of the source client.
–R *recover_target*	The name of the destination client.
–i (N,Y, R)	The recover option specifying what the Solstice Backup server should do if recovered files conflict with existing files of the same name. Here is an explanation of each option (these are not case sensitive): • **N** Skip restoring this file when there is a naming conflict. • **Y** Overwrite the existing file when there is a naming conflict. • **R** Rename this file when there is a naming conflict. The Solstice Backup server automatically appends each conflicted filename with a .R.
directory_name	The name of the initial directory in which to begin browsing.

For a complete listing of the **recover** command and all its options, refer to the **recover** man page or the *Legato Command Reference Guide*.

Example: Directed Recovery Overwriting Existing Files

If the server was named *venus*, the source client itself was *mars*, the destination client was *jupiter*, and you wanted the recovered files to overwrite any existing files with the same name, you would enter the following command at the prompt:

```
# recover -s venus -c mars -R jupiter -iY directory_name
```

Directed Recovery of Specific Files and Directories

You must specify the items you want to recover using a set of procedures to select files and directories for recovery.

To perform a directed recovery of files and directories you specified:

1. Enter the command to perform the directed recovery, using the following syntax style:

```
# recover -s server -c source_client -R destination_client -i recover_option
[directory_name]
```

2. Navigate to the directory containing file or directory you want to select for directed recovery (if you did not specify the initial directory to browse by including a *directory_name*). For example:

```
# cd /var/adm
```

3. Select the file or directory for directed recovery by entering the following command at the prompt:

```
# add file_name
```

4. Repeat this procedure for each file or directory you want to select for directed recovery.

5. View the file or directory you have selected for recovery by entering the following command at the prompt:

```
# list
```

6. Start the directed recovery by entering the following command at the prompt:

```
# recover
```

Save Set Recovery

You can use the save set recover feature to recover backed-up data that has passed the period of its browse policy, but is still in the media database. When you perform a save set recover operation, you must recover the level full backup first, then recover the other backups in level order from 1 – 9, and then recover the incremental backups.

How Save Set Recovery Works

Save set recovery allows the recovery of data from a save set rather than the regular method of selecting items from the Recover window. The data recovered can be individual files, directories, or the entire save set itself. When entire save sets are recovered—for example, an entire disk volume or partition on a disk—the recovery process is accelerated because the Solstice Backup server does not have to search for and locate each individual file in the index. Instead, the server locates the save set containing all of the files needed for recovery at once. The names of each save set for the client and the number of versions of each save set is displayed in the Save Sets dialog box.

You might want to perform a save set recovery under the following circumstances:

■ The files are no longer browsable in the Recover dialog box because their browse periods have expired or their entries have been removed with the Remove Volume command.

■ A disk crash during the backup produced an incomplete backup, making it impossible to recover the files using the recover command.

■ Using the recover command to recover data fails or is exceedingly slow (often indicative of insufficient computer memory).

Save set recovery is most efficient when restoring a large volume of data. Save set recovery is always an option for data that has passed its browse policy and is no longer accessible through the client file index. Save set recoveries avoid the overhead of the Solstice Backup file handling features.

Caution – Only root@<sbu_server> has the permission to perform a save set recovery, and, in case of a remote client, root@<sbu_server> must either be a Solstice Backup administrator, or must be entered in the Remote Access list of the client that is providing the backed-up data.

▼ To Recover Data from a Save Set

Note – Save set recovery recovers data to the machine on which the Solstice Backup Administrator program is launched. If you want to recover the data to the original client machine, you must perform this procedure on that machine.

To recover data from a save set.

1. Start the Solstice Backup Administrator program.

2. Select Save Set>Recover to open the Save Set Recover window.

3. Enter or select the client whose save set data you want to recover in the Client attribute.

   The save sets available for recovery are displayed in the Save Set attribute.

4. Select the specific save sets whose data you want to recover in the Save Set attribute. You can select as many save sets as are listed in this attribute.

   The selected save set's backup instances appear in the Instances attribute.

5. Select the particular instance of a save sets in the Instances attribute.

6. Click Details if you want to view the details of the selected save set instance, such as the backup date, number of backed-up files, save set ID number, and pool.

7. Click Recover to open the Save Set Recover Status window and customize the save set recovery (as needed).

8. Perform any or all the following:
   - Specify what save set data is to be recovered (see "To Specify What Save Set Data Is Recovered" on page 308).
   - View the volumes required to recover the specified save set data (see "To View Volumes Required for Save Set Data Recovery" on page 309).

- Relocate recovered save set data to a non-default location (see "To Relocate Recovered Save Set Data" on page 309).

- Resolve recovered save set data conflicts (see "To Resolve Recovered Save Set Data Conflicts" on page 309).

9. Click Start to begin the save set data recovery.

The recovery status appears in the Status attribute of the Save Set Recover Status window.

▼ To Specify What Save Set Data Is Recovered

You can specify what save set data is recovered while in the Save Set Recover Status window. If you want to recover the entire save set, do not change the default value displayed in the Paths to Recover attribute. When you initiate the save set recovery, the entire save set is recovered.

To limit the recovered save set data to specific files and/or directories:

1. From the Save Set Recover windows, click Recover.

1. Enter the full path of the first item you want to recover in the Paths to Recover attribute

2. Click Add.

3. Repeat Step 1 and Step 2 for each separate item you want to recover from the save set.

4. Click Start to begin the recovery.

The path you entered is then displayed in the list of save set recovery paths below the Path to Recover attribute. Once you specify individual items and their paths, the complete save set will no longer be recovered, only the particular items you entered.

When specifying directory paths for recovery, place a forward slash after a specified directory. For example, enter:

```
/usr/paul/test/
```

instead of

```
/usr/paul/test
```

This means that only a directory test will be restored and not *test1*, *test2*, and so forth. Save set recover will parse a directory name and restore all instances that match the leading string, unless limited by a slash and assuming the overwrite option in cases of a recovered data conflict.

▼ To View Volumes Required for Save Set Data Recovery

You can view a list of the required volumes necessary to recover save set data while in the Save Set Recover Status window. This window displays each volume that will need to be available to your backup device to complete the save set recovery.

To open the Save Set Recover Volumes Required window, click Volumes.

If no volumes appear, that indicates that no entries for the volume exist in the media database. Use the scanner program to re-create the client file index entries or media database entries. The scanner program can read the storage volume directly, without assistance from the Solstice Backup server. For more information, see "Recovering Client File Indexes" on page 316.

▼ To Relocate Recovered Save Set Data

You can instruct the Solstice Backup server to relocate recovered save set data from the data's original location during its backup to another place on your filesystem while in the Save Set Recover window.

To relocate recovered save set data:

1. In the Save Set Recover Status Window, click Options to open the Save Set Recover Options window.

2. Enter a path in the Relocate Recovered Data To attribute.

For more information on relocated recovered data, see "To Relocate Recovered Data" on page 297.

▼ To Resolve Recovered Save Set Data Conflicts

You can instruct the Solstice Backup server to perform a particular action in the case of a conflict between an item recovered during the save set recovery, and a like-named file already existing on the filesystem while in the Save Set Recover Status window.

To resolve recovered save set data conflicts:

1. In the Save Set Recover Status Window, Click Options to open the Save Set Recover Options window.

2. Select one of the choices in the Duplicate File Resolution attribute:

   - Rename recovered file
   - Discard recovered file
   - Overwrite existing file

For more information on recovered data conflicts and their resolution, see "Resolving Recovered Data Conflicts" on page 296.

Restoring a Save Set Entry in the Online Indexes

The following sections describe the process of restoring save set entries into one or both of the online indexes.

▼ To Rebuild Client File Index Only

If the file is not browsable (which means that the save set's browse policy has expired), but its save set is still tracked by the Solstice Backup server in the media database (which means that the save set's retention policy has not expired), follow these steps to recover the save set's entry back into the client file index:

1. Enter the mminfo command at the prompt:

   ```
   # mminfo -a -v volume_name
   ```

2. From the mminfo output, find the *ssid* with the file you want. Make sure it is not the bootstrap *ssid*.

3. Restore the save set entry in the client file index using the nsrmm and scanner programs:

   ```
   # nsrmm -S ssid -w browse_time
   ```

   Using an appropriate browse time in the future. Then enter:

   ```
   # scanner -i -S ssid
   ```

▼ To Rebuild a Client File Index and Media Database

If the save set that contains the file is not browsable and the save set is not represented in the media database, both the browse and retention policies have expired.

To rebuild the save set's entry in both the client file index and the media database:

1. Run the scanner program on the backup volume with the file you want by entering the following command at the prompt:

```
# scanner device_name
```

You might be able to determine the volume name by looking at the adhesive label on the volume. For other suggestions on finding volume names, see "To Find a Volume Name" on page 249.

2. Use the output from the scanner program to decide whether to reintroduce the contents of this volume into the client file indexes and whether the save set you want to rebuild is on this volume. (You must locate *all* the volumes that contain this save set ID.)

3. After you have determined which volumes to reintroduce into the online indexes, enter the scanner command at the prompt:

```
# scanner -i device_name
```

The scanner command prompts for a new volume until you terminate it. To rebuild the indexes completely, you must scan in *all* the volumes that contain the *ssid*.

4. Run the nwrecover program to browse the file index for the file you want to recover.

▼ To Recover an Entire Save Set to the Solstice Backup Server

To recover an entire save set directly to the Solstice Backup server's filesystem, use the **scanner** program by entering the following command at the prompt:

```
# scanner -S save_set_id device_name | /usr/lib/nsr/uasm -rv
```

This command reads all the information associated with the *ssid* from the volume and places a copy of this data on the Solstice Backup server in the exact way that it is stored on the backup volume. In other words, the backup volume may contain files for a client, but is recovered to the Solstice Backup server's filesystem.

If you want to test the command before you perform it, add the –n flag to the uasm command. The –n flag sends the output from scanner to */dev/null* and lists all the filenames contained in the save set.

You could also use rsh (or its equivalent) in conjunction with the following command to recover the save set to the client, if the save set originated on a Solstice Backup client instead of the Solstice Backup server:

```
# scanner -S ssid device_name | rsh client (cd destdir; /usr/lib/nsr/uasm
-rv)
```

▼ To Recover One File Directly from a Volume

To recover a single file from a volume, enter *one* of the following commands at the prompt:

```
# scanner -S save_set_id device_name | /usr/lib/nsr/uasm -rv filename
```

```
# scanner -S save_set_id device_name | /usr/lib/nsr/uasm -rv -m source=
dest filename
```

The –m option of uasm maps (relocates) the recovered file from the *source* to the *dest* (destination) directory.

Comparison of Recovery Methods

This section compares and contrasts the various data recovery methods provided by the Solstice Backup product.

Advantages of Index-Based Recoveries

Standard recovery and directed recovery are both based on using the client file index to locate and recover the desired data.

The following key benefits are associated with index-based recoveries:

- Directed recovery is well-suited for centralized administration. A large network can be administered for direct recovery from a single administering computer. Data recovered to an administering server can then be forwarded to a shared server drive and made available to a wide number of users.

- Directed recovery is useful if a computer is inoperable, a hard disk has crashed, or a network is no longer recognizing a computer. You can easily configure the replacement computer by recovering a set of files for the client that were backed up from another client on the network.

- Directed recovery is also convenient if you want to transfer files between two clients. For example, if one client's system file is perfect for a new client, you can recover it from the backups of the old client to the new client's hard disk.

A standard, client file index-based recovery differs from a directed recovery only in the destination of the recovered data. With a standard recovery, the Solstice Backup server recovers data to the data's original client. Thus, both kinds of client file index-based recoveries can be contrasted with a save set recovery.

You might prefer index-based recovery over save set recovery for the following reasons:

- Because users can browse the index to mark data for recovery, neither users nor administrators need to know the exact filename, location, or date of the specific files they want to recover.

  The Solstice Backup server supports multiple users (sessions) concurrently browsing a client file index, where users can expand directories to view files and subdirectories. Thus, needed files usually can be deduced from information obtained during browsing. Users are relieved of the burden of remembering explicit filenames from months past.

- Users can query the client file index with a target time for which they want to rebuild the file hierarchy. The Solstice Backup server constructs for the user a view of the filesystem that was written closest to the browse target time without exceeding it.
- The Solstice Backup server automatically re-creates the data exactly as it existed at the time that the backup occurred.

  For example, if a user requests recovery of a backup version that is a level 5 backup, the Solstice Backup server first recovers the full backup on which that level 5 depends, then applies the level 5 backup. Any renaming of files or deletions that occurred between the full and level 5 backups are handled as well.

Disadvantages of Index-Based Recoveries

The disadvantages of index-based recovery are situation-specific:

- If the data a user wants to recover has passed its browse policy period, its entries have already been removed from the index. Entries for the files can be recovered and restored to the client file index using the scanner command, but this can be time-consuming or inconvenient. For more information on scanner, see "Recovering Client File Indexes" on page 316.
- An index-based recovery takes time to mark each file for recovery.
- An index-based recovery requires memory proportional to the number of files marked for recovery.

For all these reasons, if recovery is planned for a significant number of files, you might prefer to use a save set recovery.

Disadvantages of Save Set Recoveries

With save set recovery:

- You must locate the full save set and its relevant level saves and incrementals, and restore all of those save sets in the correct order.
- The recovery can restore more data than was requested. Files that were deleted after the last full save set backup are also recovered.

A save set recovery follows the basic steps of an index-based recovery, but without the advantage of client file index handling that the Solstice Backup server provides automatically with index-based recovery.

⚠ **Caution –** For this reason, if you are recovering a number of save sets that were backed up at different levels, you must be careful to execute the save set recovery requests in the correct order, beginning with the full backup. For example, if file X

existed at time A, but was deleted prior to time B, file X will be recovered during a save set recover at time B. Solstice Backup file handling is not available for save set recovery, so file X remains on your Solstice Backup server during the recovery operation.

It is possible that disk space could become scarce if too many previously deleted files or directories are recovered. If the save sets you recover are full rather than incremental, running out of disk space will be far less of an issue because you will not recover multiple copies of the same files or directories.

Typically, during a recovery you want to force the **nwrecover** program to overwrite existing files. Overwriting files is even more important with save set recovery, because the same file might be recovered multiple times, with each successive version of the file coming from a later save set.

Deciding between Index-Based and Save Set Recovery

This section summarizes the reasons for using either index-based or save set- based recovery.

- Client file index-based recovery is most appropriate under the following conditions:
  - You cannot determine which save sets to recover.
  - Recovering extra files is not desired.
  - Using save set recovery, you may have to perform multiple recoveries to restore the file to the correct state. For example, if a file is first backed up as full, then backed up at an incremental level, you will have to restore two save sets—first the full backup, then the incremental backup—to restore the file to the state it was in after the incremental backup. Using index-based recovery, only one recovery is necessary. The Solstice Backup server handles the incremental backup issue internally, prompting you, if necessary, for the appropriate volumes.
- Save set recovery is most appropriate under the following conditions:
  - You can determine the correct save sets to recover.
  - There are only a few save sets to recover for each filesystem.
  - Recovering extra files is acceptable.
  - Free disk space is not a problem.
  - Your most recent save sets are from full, not incremental backups.
  - The browse policy has expired, and you do not want to use the scanner program to re-create the client file index entries.

Recovering Client File Indexes

The following sections explain different procedures for recovering the online index.

▼ To Recover Client File Indexes

Unlike regular backed-up data, recovery of client file index information is performed using the nsrck program with the -L7 option to perform a level sever consistency check. However, a level 7 check merges recovered data into the existing client file index; therefore, if the existing client file index has become corrupt, you must first purge the corrupt save set information from the index using a level 5 consistency check before recovering the index from a backup volume.

To recover a client file index of a client:

1. Run the nsrck command with the -L5 option to purge corrupted save set information from the index:

   ```
   # nsrck -L5 client_name
   ```

2. Run the nsrck command with the -L7 option to recover the index:

   ```
   # nsrck -L7 client_name
   ```

   You can also completely remove a client file index by using nsrck -R -Y *client_name*, then use nsrck -L7 *client_name* to recover the index.

▼ To Recover Client File Indexes for a Specific Time and Date

You can recover a client file index for a certain time and date, using the -t "time date" option with nsrck. The time and date formats are in nsr_getdate format. For details on the nsrck and nsr_getdate commands, refer to the *Legato Command Reference Guide* or the appropriate man pages.

Client file index recovery using the Solstice Backup product is nondestructive, so when you recover the client file index as of a time and date in the past, nsrck will add the entire contents of the index as of that time and date, to the current contents of the index. Generally, this means your client file index will become much larger.

The benefit of recovering the client file index as of a specific time and date is that you can browse save sets that had passed their browse policy but were still recoverable. The save sets referred to by the recovered client file index will be marked as browsable. They will remain browsable for the length of time they were originally browsable.

To recover a client index for a specific date, enter the following command at the prompt:

```
# nsrck -t "time mm/dd/yyyy" -L7 client_name
```

Example: Recovering a Client File Index for a Specific Time and Date

In this example, you are recovering a client file index to 4 PM, March 2 2001:

```
# nsrck -t "04 pm 03/02/2001" -L7 client_name
```

Relocating Recovered Online Indexes

You can relocate recovered online indexes depending on the release of Solstice Backup software that performed the backups.

Release 6.0 or later Software

You can recover online indexes to a different directory only if all the backups were made using the Solstice Backup release 6.0 or later software.

Release 5.5 and Earlier Software

You cannot relocate recovered online indexes to a different directory if they contain any backups performed using the Solstice Backup release 5.5 or earlier software. Instead, you must recover them to the directory where they originally resided. After you recover the indexes to their original location, you can move them to another directory if necessary. If you attempt to recover the indexes to a different directory, you will receive a warning message stating that the recovery for the client index was unsuccessful.

The scanner Program

The scanner program reads the volume directly, without accessing information that might be available in the media database. If no entries exist in the media database, you must take an educated guess at which volume contains the data to recover.

The scanner program is used to import data that has passed its browse and retention time, or was backed up by a different Solstice Backup server. You should not use scanner to restore the bootstrap in the event of a disaster in which your media database or server's res files are lost. Instead, you should use mmrecov. You can use scanner -B to locate bootstrap save sets on a volume, and then use mmrecov to restore the bootstrap. For more information, refer to the *Legato NetWorker Disaster Recovery Guide*.

Mount the volume and execute the **scanner** command with any of its applicable options. The option that is most appropriate depends entirely upon the situation, your level of knowledge about the location of the data, and the nature of the recovery.

The following are the most commonly-used **scanner** options:

TABLE 7-3   Scanner Options *(1 of 2)*

Option	Meaning
-m *device*	Rebuild the media database entries for everything on the storage volume.
-i *device*	Rebuild the media database and client file index entries for everything on the storage volume.
-S *save_set_id* -m *device*	Rebuild the media database entries for a specific save set on the storage volume.
-S *save_set_id* -i *device*	Rebuild the media database and client file index entries for a specific save set on the storage volume.

TABLE 7-3   Scanner Options *(2 of 2)*

Option	Meaning
−c *client_name*	Process only save sets that come from the specified Solstice Backup client machine.
−N *save_set_name*	Processes only save sets specified by *save_set_name*.
−S *save_set_id*	Sends a specific save set to standard output.
−x	Introduces further processing commands on a per-save set basis.
−B	Locates the bootstrap save sets on the volume

Note – For information about using scanner to import save set information from a volume backed up by a different Solstice Backup server, see "Importing Save Set Information with the scanner Program" on page 319.

Detailed information about scanner program syntax and examples of how scanner options can be customized to meet the needs of a particular recovery situation are available in the *Legato Command Reference Guide* on the Documentation Suite CD-ROM.

The entries placed in the client file index receive the original browse and retention policies as of the time you scan in the save set. If a save set originally had a browse time of one month and retention time of three months, and that save set had expired, and you scanned that save set back in, it would remain browsable for 1 month and recoverable for three months after you scanned it in.

The entries placed in the media database also receive the original browse and retention policies as of the time you scan in the save set, and the volume mode is listed as "read-only." For more information on managing the client file index and media database entries, see "Storage Management Operations" on page 240.

If you upgrade a Solstice Backup server to 6.0 or later, the current browse and retention policies are applied to save sets that had been scanned into the client file index and the media database.

Importing Save Set Information with the scanner Program

The scanner -i program can be used to import client file index entries for save sets on a volume that was backed up by a different Solstice Backup server. However, with Solstice Backup release 6.0 and later, each client is associated with a unique

Client ID. During backup, the Client ID is recorded with each save set. When clients with the same name are configured on different Solstice Backup servers, they will have different Client IDs.

The scanner program uses the Client ID, rather than the client name, to identify clients. As a result, even when imported save sets belong to a client that has the same name on both servers, scanner is not able to associate imported data with the appropriate client.

To import the client data in this case, you must create a new client on the new server based on the Client ID from the old server. You cannot simply delete the old client and recreate it with the new Client ID. The Solstice Backup server retains the old Client ID even after the client is deleted, so that all save sets indexed for the old client can be recovered. When a new client is created using the same name as a deleted client, the new client will automatically have the Client ID of the deleted client. You can either create a new client, or you can rename the existing client by following the procedure found at "To Change a Client's Name" on page 145 and then create a new client using the Client ID from the old server.

▼ To Create a New Client Based on the Client ID

To create a new client based on the Client ID:

1. Determine the Client ID on the old server. For information about determining the Client ID, see "The Client ID" on page 144.

2. Start the Solstice Backup Administrator program and select Clients>Client Setup to open the Clients window.

3. Click Create.

4. In the Name attribute, enter a name for the client. This can be the same name that was used on the old server, but it cannot be the same name as a client that already exists on the new server.

   If a client with the appropriate name already exists on the new server, use the following format for the name of the client:

   ```
   ~hostname-#
   ```

   Where *hostname* is the hostname of the client. For example, if the client's hostname is jupiter, and a client named jupiter already exists on the new server, you would enter the following for the client name:

   ```
   ~jupiter-1
   ```

5. In the Client ID attribute, enter the Client ID determined in Step 1.

6. Complete other attributes as necessary, and click Apply.

Once the new client has been created, import the client data using `scanner -i`. You can now perform directed recoveries to restore the client data. For more information about directed recoveries, see "Directed Recovery" on page 299.

Note – If `scanner -i` or `scanner -m` is used to import data before the Client resource has been configured on the new server, the Client ID for the imported save sets is maintained in the media database. If a client of the same name already exists on the new server, `scanner` will store the client name in the format described in Step 4. You can then create the client based on the Client ID, using the above procedure. The entries in the media database will be associated with the correct client. However, you must run `scanner -i` again after creating the Client resource in order to import save set information into the client file index.

Data Archiving

This chapter describes the features and functionality of the optional NetWorker archive feature.

The main sections in this chapter are:

Overview

The NetWorker archive feature provides archive services to Solstice Backup client computers. The archive process captures files or directories as they exist at a specific time, and writes the data to special archive storage volumes, which are not automatically recycled. After the archive process completes, you can delete ("groom") the original files from the disk to conserve space.

Because the Solstice Backup server treats archive data differently than backup save set data, archive data must be written to separate pools.

Some of the benefits of using the NetWorker archive feature include:

- Files that have been archived can be deleted from your primary disk storage to make space for newer files, and retained in archive volumes for quick retrieval in the future.
- Archived data is never subject to automatic recycling, so it cannot be accidentally overwritten.

- Archived volumes that use the PC-style archiving method maintain their entries in the media database indefinitely. (PC-style archiving is only available for PC clients, such as Microsoft Windows NT or Microsoft Windows 2000).

Archives Compared to Backups

Archive save sets are similar to backup save sets; the principal difference is that archive save sets have no expiration date. Also, archives are always full; there are no levels 1–9 or incremental archives.

Browse and retention policies do not apply to archive data. Having no retention policy ensures that archive data is never eligible for automatic recycling.

Requirements

Before you begin licensing and using the NetWorker archive feature, make sure that you have the following:

- A device, either stand-alone or in an autochanger or silo, connected to a Solstice Backup server or storage node. If you are cloning your archives, you must have at least two devices available.
- A temporary or permanent enabler code to license the product after any evaluation period is over.

Evaluating Solstice Backup Archive

The NetWorker archive feature is an optional module that you must purchase and license separately from other Solstice Backup software components. The Solstice Backup client and server installations include the files necessary to use NetWorker archive feature. However, the functionality only works when the archive feature is properly licensed in the Solstice Backup server.

The evaluation conditions and procedures are the same for the archive feature as for other Solstice Backup products. For more information on licensing procedures, see the Solstice Backup Roadmap.

Archiving Permissions

The following sections describe permissions required to use the Archive feature.

Enabling Archive Services for the Client

After you enter the archive enabler code for the Solstice Backup server, all clients of that server are enabled for the NetWorker archive feature by default. To disable the Archive feature for a specific client, set the Archive Services attribute in the Client resource to Disabled.

To archive data that resides on the NetWorker server, make sure that the Archive Services attribute is enabled in the Clients resource for the server. For instructions, see .

Caution – When you enable the Archive Services attribute for a Client resource, you also enable the Archive Services attribute for all other clients of the same name on that server. For example, if the Legato NetWorker Module for a database application and the Solstice Backup client software are installed on the same computer and both back up to the same Solstice Backup server, both Client resources have the same name. The Archive Services attribute is either enabled for both or disabled for both.

Enabling or Restricting Archive Access

The Archive Users attribute in the Client resource specifies the users who are allowed to archive and retrieve data. If no users are listed in this attribute, only Solstice Backup administrators with read permissions to the data can archive or retrieve the data. To enable users who are not Solstice Backup administrators to archive and retrieve data, include their user names in this attribute. They must have read permissions to the data in order to archive or retrieve the data.

By default, only the following users can retrieve archived data:

- Solstice Backup administrators
- The owners of the archived files
- The system administrator (`root@client` on UNIX systems or a user with Administrator privileges on Microsoft Windows systems)

If other users need to retrieve data they do not own, public archives access should be enabled.

Enabling Public Archive Access

To allow users listed in the Archive User attribute of the Client resource to retrieve any archived file from that client, enable the Public Archives attribute in the Server resource. If, during recovery, your operating system allows you to change the ownership of archived data to that of the original owners, then the retrieved files display their original ownerships. Otherwise, the user who retrieves the files becomes the owner of the files.

You cannot omit specific clients from the retrieval permissions when you enable the Public Archives attribute.

How the Solstice Backup Server Archives Data

The Solstice Backup software provides two preconfigured pools to receive archived data: the Archive pool and the PC Archive pool. You can also create custom pools for archiving by selecting Archive for the Pool Type attribute in the Pool resource. During the archive operation, the Solstice Backup server writes data to storage volumes that belong to an Archive-type pool. Information about archive data is tracked in the Solstice Backup server's media database.

If you use the preconfigured PC Archive pool, or if you create a custom Archive pool that has the Store Index Entries attribute in the Pool resource set to Yes, information about individual files in the archive save set are tracked in the client file index. The client file index entries that are generated during an archive are backed up to volumes from the Default pool during the next scheduled backup. Index entries are not generated when the Store Index Entries attribute in the Pool resource is set to No.

The Solstice Backup server tracks the volumes used for archives separately from those used for backups. You cannot archive files to a backup volume, or back up files to an archive volume. An archive volume must be loaded and mounted in the server device to complete an archive.

Whether you initiate the archive on the client or the server, the archive is performed by the client's archive program (nsrarchive), which is initiated by the client's nsrexecd service. The nsrarchive program performs manual archives from the client's command line. You can schedule archives from the server or client using the Archive Requests resource in the Solstice Backup Administrator program.

FIGURE 8-1 on page 327 illustrates how the Solstice Backup software archives data.
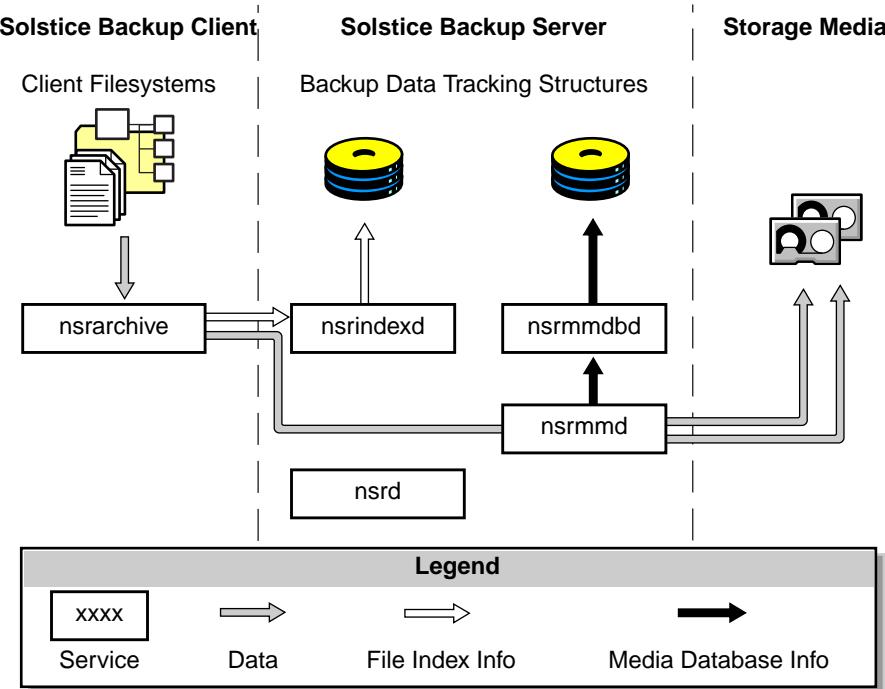
FIGURE 8-1   Overview of Archive Operation

UNIX-Style Compared to PC-Style Archiving

The Solstice Backup server supports two styles of archiving: UNIX style (nonindexed) and PC style (indexed).

UNIX-Style (Nonindexed) Archiving

When data is archived using UNIX-style archiving, entries are not added to the client file index. When this data is retrieved, the entire save set is retrieved and you cannot browse or recover individual files. For UNIX clients, this is the only style of archiving that is supported. PC-based Solstice Backup clients, such as those for Microsoft Windows, can use either UNIX- or PC-style archiving

PC-Style (Indexed) Archiving

Data archived using PC-style archiving has entries added to the client file index. As a result, you can browse and select individual files for recovery.

PC clients that want to take advantage of indexed archiving must either use the preconfigured PC Archive pool or enable the Store Index Entries attribute of the Pool resource associated with the archive volume.

The client file index entries that are generated during an archive are backed up to volumes from the Default pool during the next scheduled backup.

Note – Although you can select the PC Archive pool from a UNIX or a PC client, the ability to browse and recover individual files is only supported for PC clients using the PC-based Solstice Backup User program (`winworkr.exe`). Recovering individual files from archived save sets is not supported on Solstice Backup clients in UNIX environments.

About Archive Pools

The Solstice Backup software provides two preconfigured pools to receive archived data: the Archive pool and the PC Archive pool. You cannot change the settings for these preconfigured pools, although you can create custom pools for archiving data. Custom pools can use either UNIX-style or PC-style archiving.

Preconfigured Archive Pool

Also called UNIX-style archiving, the preconfigured Archive pool does not have a browsable client file index associated with it. Therefore, you cannot retrieve individual files from the archive save set. You must retrieve the entire save set.

Preconfigured PC Archive Pool

Also called PC-style archiving, the preconfigured PC Archive pool stores entries for individual files in the client file index. Although you can select the PC Archive pool for both UNIX and PC clients, you can only browse or recover client file index entries from a PC client.

Caution – Using the PC Archive pool might create a large client file index that never expires. To conserve disk space, purge or compress the index.

Creating Custom Archive Pools

Two attributes in the Pool resource are used to distinguish Archive Pools from other pools:

- The Pool Type attribute must be set to Archive. This attribute tells the Solstice Backup server that volumes belonging to this pool are used for archiving.
- The Store Index Entries attribute determines whether the archive is a UNIX-style or PC-style archive:
  - If this attribute is set to No, entries are not written to the client file index (UNIX-style archiving).
  - If this attribute is set to Yes, entries are written to the client file index (PC-style archiving). This method is only supported on PC clients.

For details and procedures about creating pools, see "Pools" on page 194.

Archiving Data

You can request manual archives from the client, or you can schedule archives from the server.

This section outlines the following procedures:

▼ To Enable Archive Services for a Solstice Backup Client

To enable the archive services for a Solstice Backup client:

1. In the Solstice Backup Administrator program, select Client>Client Setup to open the Clients window.

2. To create a new client, click Create and complete the client information. To edit an existing client, select the client from the list provided.

3. Select Enabled from the Archive Services attribute to enable archive services for the client.

4. Enter all appropriate users in the Archive Users attribute. For details about setting this attribute, see "Enabling or Restricting Archive Access" on page 325.

5. Click Apply.

Manually Archiving Data

You can archive data at any time using the manual data archiving method. Manually archiving data is similar to manually backing up data.

You can designate the entire filesystem, specific directories, or individual files for manual archiving. Note that if you select a directory for archiving, all of its nested subdirectories and files are also selected. If you want to exclude certain files from manual archiving, you must deselect those items using the nwarchive program.

▼ To Request a Manual Archive from a Solstice Backup Client

To request a manual archive:

1. Start the nwarchive program as follows:

   – To archive to the default Solstice Backup server, enter the following command:

     # **nwarchive &**

   – To archive to a Solstice Backup server other than your default server, enter the following command:

     # **nwarchive -s** *server* **&**

   The nwarchive program displays a representation of the client's filesystem, similar to the nwbackup program used for manual backups.

2. Select Mark>Mark or click the Mark button to mark each file or directory for archiving. When you mark an item for archiving, a check mark appears next to that item.

   To unselect an item currently marked for backup, select Mark>Unmark or click the Unmark button.

3. Select File>Start Archive or click the Start toolbar button to begin the archive.

4. Enter a comment in the Annotation attribute. This annotation is used to identify each archive save set during retrieval.

5. Select the appropriate archive pool from the Archive Pool attribute. Only pools with their Pool Type attribute set to Archive will be listed.

6. Select the appropriate response for the following data archiving attributes:

- Compress: To compress the data as it is archived, select Yes for the Compress attribute. Compression and decompression of data typically uses more of your server resources and may slow data archiving and retrieval.

- Clone: To write a copy of each archive save set to a volume from an Archive Clone pool, select Yes for the Clone attribute.

  If you enable cloning, enter or select an archive clone pool in the Clone Pool attribute. Only pools with their Pool Type attribute set to Archive Clone will be listed.

- Verify: To instruct the Solstice Backup server to check the integrity of the data on the storage volume, select Yes for the Verify attribute.

- Grooming: To instruct the Solstice Backup server to remove the archived files from your disk, select Yes for the Grooming attribute.

7. Click Archive.

The Solstice Backup server displays the Archive Status window, which monitors the progress of the archive. When the Solstice Backup server is finished archiving, you see a message similar to the following in the Archive Status dialog box:

```
Archive completion time: 2-21-95 5:18p
```

Using nsrarchive to Manually Archive Data

You can also use the command-line program nsrarchive to manually archive data. For more information, see the man page for this command, or refer to the *Legato Command Reference Guide*.

Scheduled Archives

Scheduled archives are performed by the Solstice Backup server and allow you to schedule the archiving of client data using the Solstice Backup Administrator program. Unlike scheduled backups, scheduled archives run only once. The advantage of a scheduled archive is that you can run the archive when network traffic and computer use are low.

▼ To Schedule an Archive on the Solstice Backup Server

Note – Before you can schedule an archive request, the Archive Services attribute in the Client resource must be enabled. For details, see "To Enable Archive Services for a Solstice Backup Client" on page 329.

To schedule an archive:

1. In the Solstice Backup Administrator program, select Customize>Archive Requests to open the Archive Requests window.

2. Click Create.

3. For the Name attribute, enter a name for the archive request (limit of 64 characters).

4. For the Annotation attribute, enter a comment. This annotation is used to identify each archive save set during retrieval.

5. Select one of the following options in the Status attribute:
   - Start Now: Starts the archive as soon as you click Apply in the Archive Request window.
   - Start Later: Starts the archive at the time you specify in the Start Time attribute. Enter a time for the archive request to start in the Start Time attribute. Enter the time value in 24-hour time format, such as 23:59 for 11:59 p.m. The default start time value is 3:33.

6. For the Client attribute, enter the archive client's hostname. If you want the Solstice Backup server to archive its own files, enter the server's hostname.

7. For the Save Set attribute, enter the pathnames of the files and/or directories you want to archive.

Note – If you want to create an archive request for all of the client's save sets, enter the value All. If you enter this value, be sure to select None for the Grooming attribute. If you select Remove for the Grooming attribute, the Solstice Backup server will delete all of the client's archived save sets.

---

8. For the Directive attribute, enter or select a directive. For more information on directives, see "Directives" on page 178.

9. From the Archive Pool list, select the appropriate pool.

10. Select the appropriate response for the following data archiving attributes:
   - Verify: To instruct the Solstice Backup server to check the integrity of the data on the storage volume, select Yes for the Verify attribute.
   - Clone: To write a copy of each archive save set to a volume from an Archive Clone pool, select Yes for the Clone attribute.

     If you enable cloning, enter or select an archive clone pool in the Archive Clone Pool attribute.
   - Grooming: To instruct the Solstice Backup server to remove the archived files from your disk, select Yes for the Grooming attribute.

11. For the Archive Completion attribute, enter a notification for the Solstice Backup server to use after completing the data archive. For more information on notifications, see "Notifications" on page 99.

12. Click Apply.

To view information about the status of an archive request, display the Archive Request Details dialog box. See "To View Details of a Scheduled Archive" on page 337.

▼ To Edit an Archive Request

To edit an archive request:

1. In the Solstice Backup Administrator program, select Customize>Archive Requests to open the Archive Requests window.

2. Select the archive request from the list displayed in the Archive Requests attribute.

3. Edit the attributes of the archive request and click Apply.

Note – You cannot edit an archive request that is currently in use.

▼ To Delete an Archive Request

To delete an archive request:

1. In the Solstice Backup Administrator program, select Customize>Archive Requests to open the Archive Requests window.

2. Select the archive request from the list displayed in the Archive Requests attribute.

3. Click Delete.

4. When prompted to confirm that you really want to delete that archive request, click OK.

Retrieving Archived Data

This section describes how to use the NetWorker archive feature to retrieve archived data.

Retrieval Permissions

In order to retrieve archived data, the following restrictions apply:

- You must have read permissions to the archived data.
- If the server's Public Archives attribute is enabled, all users listed in the client's Archive Users attribute can retrieve data (as long as they have read permissions to the data).
- If the server's Public Archives attribute is disabled, only the Solstice Backup administrator or the owner of the file can retrieve the data.

For more information, see .

Note – If, during recovery, your operating system allows you to change the ownership of archived data to that of the original owners, then the retrieved files display their original ownership. Otherwise the user who retrieves the files becomes the owner of the files.

▼ To Retrieve UNIX-Style (Nonindexed) Archived Data

Data archived using UNIX-style archiving does not maintain information about the archived save set in the client file index. Therefore, data is retrieved for an entire save set.

To retrieve nonindexed archived data:

1. Mount the archive volume in the appropriate storage device.

2. Start the `nwretrieve` program.

3. To limit the archived save sets found during the query, complete one or more of the following attributes:

   - Client: Enter the name of the client whose archived data you want to retrieve.

- Save set name: Enter the name of the archived save set you want to retrieve.

- Search annotation for: Enter all or part of the annotation assigned to the save set when it was archived. To aid the search, use wild-card symbols in this attribute.

4. Click Query to search for the save sets that match the criteria specified in Step 3. Matching archived save sets are displayed in the Archives window.

   You can view the details of an archived save set, such as its size, save set ID, and archive date, by selecting the archived save set and clicking Details.

5. To begin the retrieval, select the archived save set you want to retrieve and select File>Start Retrieve or click the Start toolbar button. The Retrieve Status window appears.

   If you want to relocate the files to a different location, enter the new location in the Relocate To attribute of the Retrieve Status window. To automatically overwrite existing files, enable the Overwrite Existing Files attribute.

6. Click OK.

   The retrieval status is displayed in the Status attribute of the Retrieve Status window. You may be prompted to indicate an appropriate action for the Solstice Backup server to perform when it encounters files with duplicate filenames during retrieval.

Using Save Set Recovery to Retrieve UNIX-Style (Nonindexed) Archives

You can also recover archived save sets using save set recovery. For information, see "Save Set Recovery" on page 306.

Recovering PC-Style (Indexed) Archived Data

Data archived using PC-style archiving maintains information in the client file index about individual files in the archive. However, you cannot browse and select archived indexed data from a UNIX client. To recover indexed archive data for Microsoft Windows clients, you must use the PC-based Solstice Backup User program (winworkr.exe), which is the equivalent of the UNIX nwarchive and nwretrieve programs.

Archive Request Management

This section describes how to accomplish basic tasks involving Solstice Backup archive requests.

▼ To Start a Scheduled Archive at Any Time

When you want to start a scheduled archive right away, rather than waiting for the scheduled time, the Solstice Backup server overrides the archive request's scheduled start time and immediately archives the data specified in the archive request.

To start a scheduled archive:

1. In the Solstice Backup Administrator program, select Server>Archive Request Control to open the Archive Requests window.

2. Select the name of the archive request in the Archive attribute.

3. Click Start.

▼ To Stop a Scheduled Archive While in Progress

To stop an archive request in progress:

1. In the Solstice Backup Administrator program, select Server>Archive Request Control to open the Archive Requests window.

2. Select the name of the archive request you want to stop in the Archive attribute.

3. Click Stop.

When you stop a scheduled archive in progress, the archiving action stops immediately and the archive request is listed as stopped in the Archive Requests attribute of the Archive Requests Control window. To restart the archive request, see "To Start a Scheduled Archive at Any Time" on page 336.

▼ To Disable a Scheduled Archive

To disable an archive request:

1. In the Solstice Backup Administrator program, select Server>Archive Request Control to open the Archive Requests window.

2. Select the name of the archive request you want to disable in the Archive attribute.

3. Click Disable.

▼ To View Details of a Scheduled Archive

The Archive Request Details dialog box provides information about the completion of an archive request:

- The Completion Time displays the time the archive finished. The time it took to complete the archive is the difference between the completion and start times of the archive.
- The success of the archive request is either completed, failed, or partial.

To view Archive Request Details:

1. In the Solstice Backup Administrator program, select Server>Archive Request Control to open the Archive Requests window.

2. Select the name of the archive request whose details you want to view in the Archive attribute.

3. Click Details to open the Archive Request Details window.

Cloning

The Solstice Backup server can make duplicates, or *clones*, of complete volumes of stored data as well as individual save sets. Cloning allows you to make duplicates of either complete volumes or individual save sets for the purpose of secure offsite storage, easy transfer of data from one location to another, or verification of backups.

The main sections in this chapter are:

- "Volume Cloning" on page 339
- "Save Set Cloning" on page 345
- "Other Information and Procedures" on page 356

To perform cloning operations, you must be a Solstice Backup administrator.

Volume Cloning

Volume cloning is the process of reproducing complete save sets from a storage volume to a clone volume. You can clone save set data from backup, archive, or migration volumes. You can clone save sets automatically, as part of any of these operations, or manually at another time.

How Volume Cloning Works

Volume cloning is a two-part operation:

1. The Solstice Backup server recovers data from the source volume.

2. The Solstice Backup server writes the data to a clone volume (a volume that belongs to a *clone pool*).

Caution – Cloning requires at least two active devices: one for reading the source volume, and one for writing the new, cloned data to another volume. This means there must be at least two devices enabled on the server in order for cloning to succeed.

During cloning, data is reproduced from the source volume to the clone volume. Cloning does not involve data stored on the clients or server. The Solstice Backup server allows only one clone of a save set per volume. Therefore, if you specify three clones of a save set, each clone is written to a separate volume.

When you clone data, different capacities of storage media may mean that more or fewer clone volumes are required than were required for the original data.

When you specify that a particular volume is to be cloned, the Solstice Backup software uses the save sets on the specified volume as the source data. When you specify that a particular save set is to be cloned, the Solstice Backup software first determines whether that save set already has a clone. If multiple clones of a save set exist, clones of save sets on volumes in a library are generally selected as the source data, rather than a volume that is manually mounted. By using the command line, you can specify the precise save set clone to use as the source.

The Solstice Backup Administrator program displays cloning information in the Messages window. Clone-related messages are also logged to the Solstice Backup message file `/nsr/logs/messages`.

Volume Cloning Requirements

The following are required for cloning volumes with the Solstice Backup server:

- Multiple storage devices, minimum of two: one to read the existing data and one to write the cloned data.
- The *destination* volume must be a different volume from the *source* volume, and the destination volume must belong to a clone pool.

  The source volume is the original volume and the destination volume is the volume you clone the data to.

Automatic Compared to Manual Cloning

Automatic volume cloning (cloning associated with a scheduled backup) is performed after all backup operations are complete. The savegroup completion report that is issued after a scheduled backup also includes a report of the success or failure of the cloning operation for each save set.

The location of the devices where the clone data is written is established in the Client resource's Clone Storage Nodes attribute. You can add or remove the names of storage nodes and the Solstice Backup server at any time.

If you want to perform cloning separately from a scheduled group backup operation, you must perform the cloning manually, volume by volume, or from the command line. If you execute cloning manually, no report is generated.

To manually clone volumes, you must be a Solstice Backup administrator. Verify that you are included in the Administrator attribute of the Server resource.

Volume Cloning Performance

In general, a volume write occurs at the same speed whether it is part of a backup operation or a cloning operation. However, if a cloning operation is automatically requested as part of a scheduled backup, you may experience a performance degradation in other scheduled backups that follow.

The Solstice Backup server usually attempts to complete one group's scheduled backup before initiating a scheduled backup for another group. However, the Solstice Backup server considers a group backup is finished when the backup operations are complete, not when any automatic cloning is complete. Therefore, if another group starts its backup while the previous group's clone operation is under way, there may be a contention for nsrmmd resources or specific volumes.

To avoid this problem, you may decide to refrain from automatic cloning and instead initiate a single cloning operation using the nsrclone command as part of a scheduled task that runs at a nonpeak time after all backups are complete. For more information, refer to the nsrclone man page or the nsrclone entry in the *Legato Command Reference Guide*.

Cloning Compared to Duplication of Volumes

When a volume is cloned, it is the save sets contained on the volume that are actually cloned. As a result, it is possible that more than one volume is required for a single volume clone, because the last save set on the volume may continue onto another volume. When this is the case, the save set is still cloned and an additional volume is required to complete the clone.

Also, the cloning operation leaves traceable information entries in both the client file index and the media database. Being able to track cloned data is an important distinguishing feature between cloning and an operating system or hardware device copy operation.

Problems with Volume Duplication Using a Copy Operation

If you rely on an operating system or hardware device copy operation to duplicate volumes, you must first ensure that the destination can hold all the save sets contained on the source volume. If the last save set continues onto another volume, this save set is not automatically duplicated onto a second volume (as they are with cloning). This could result in lost data.

In addition, Solstice Backup software does not recognize a duplicated volume, because it does not have a valid Solstice Backup label. If auto media management is enabled and the duplicated volume is left in the autochanger, the unlabeled volume is considered eligible for labeling and reuse during a scheduled backup. For more information on auto media management, see "Auto Media Management" on page 235.

Similarly, it is possible to make an exact copy of a Solstice Backup archive volume. However, the annotation information associated with each archive save set is stored in the Solstice Backup server's media database, not on the archive volume itself. Therefore, a duplicate of an archived save set does not include this annotation information. Furthermore, if the entry for the original archive save set is removed from the media database, the annotation information that describes it is also removed.

Volume Cloning and Data Tracking Information

A cloning operation does not insert entries into the client file index. Cloned save sets are only tracked through the media database. During cloning, the location of a cloned save set is added to the existing save set entry in the media database. That is, each save set clone shares the same *save set ID* (*ssid*) as the source save set. All characteristics that are true for the source save set are also true for the clone save set:

■ If the source save sets are still browsable, the clone status is also browsable.

■ If the source save sets have passed their browse policies, the clone status is recoverable.

Volumes that belong to a clone pool are also tracked through volume entries in the media database. The fact that all save sets share the same media database save set entry has implications for the following actions, which are executed on a *per-save-set* basis and not on a *per-volume* basis:

■ Changing the mode of a cloned volume (of save sets)

■ Purging a volume (of save sets) from the client file index

■ Deleting a volume (of save set locations) from the media database

Changing the Mode of a Volume Clone to Recyclable

In some cases, you might want to manually change the mode of a clone volume to *recyclable*, with the intent of reusing that clone volume. However, keep in mind that the mode of a volume only changes to recyclable when all the save sets on that volume are recyclable. Therefore, when you change the mode of the clone volume to recyclable, you effectively change the status of all save sets on the volume to recyclable. Because the save sets share the same entry in the media database, there is no distinction between original and clone save sets. The end result is that all the save sets that reside on the now-recyclable volume, or on any other volume, become candidates for immediate recycling.

If you want to reuse a particular clone volume and still protect the save sets that exist on other volumes from inadvertent data loss, first change the mode of the volumes that you want to protect to manually recycle. This way, the Solstice Backup software cannot automatically recycle the volume. Then, you can safely change the volume that you intend for reuse to recyclable mode.

Purging a Clone Volume

Similarly, if you *purge* a clone volume, you effectively remove from the client file index all the file entries associated with all the save sets that reside on the clone volume (in whole or in part).

Deleting a Clone Volume

If you *delete* a clone volume, the nsrim index management program locates the entry in the media database for each save set that resides on the clone volume. The nsrim program then marks the information about the location of one of the save set clones for deletion from the entry. This action is performed for each save set entry. In addition, nsrim marks the entry for the particular clone volume (identified by its volume ID number) for deletion from the database.

Volume Cloning and Recovery

A clone volume is used for recovery any time the Solstice Backup server attempts to recover a particular save set and one of the following has occurred:

- The original save set volume has been deleted.
- The status of the original save set has been marked *suspect*.

    For more information about changing the status of a save set, see "To Change the Status of a Save Set" on page 155.

You can execute the `scanner` program on a clone volume to rebuild entries in the client file index, the media database, or both. After you re-create the entries, traditional recovery is available. For more information on the `scanner` program, see "The scanner Program" on page 318. For information on how to recover data with the `scanner` program, refer to the *Legato NetWorker Disaster Recovery Guide*.

Creating Clone Volumes

You can manually clone all of the data on a volume in the Clone Volume dialog box, as shown in FIGURE 9-1 on page 344. This dialog box provides information about the volumes for the server currently in use, and is a starting point for the cloning process.
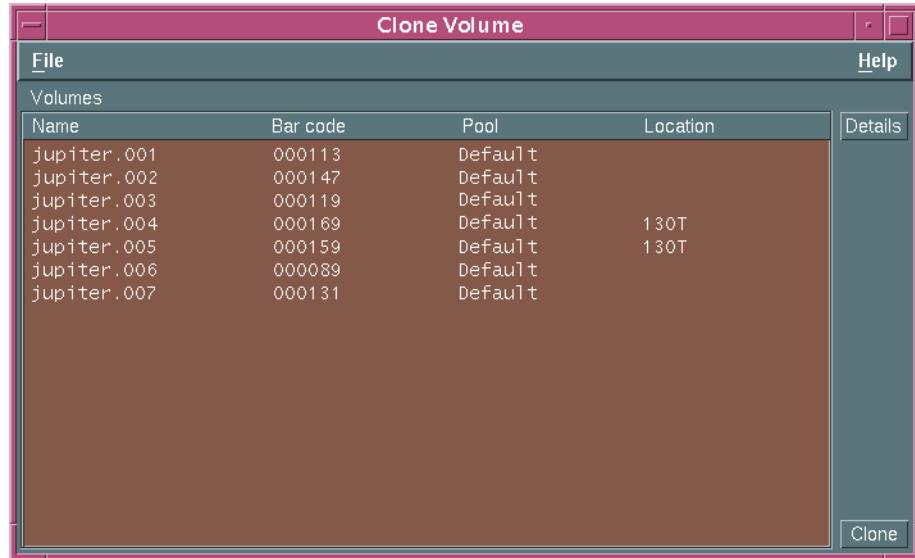
FIGURE 9-1   Clone Volume Dialog Box

▼ To Create a Clone Volume

To create a clone volume:

1. In the Solstice Backup Administrator program, select Media>Clone Volume to open the Clone Volume dialog box.

2. Select the volume to be cloned from the Volumes list.

3. Click Clone to open the Clone Volume Status window.

4. Enter the name of the clone pool to which the clone volume will belong. Cloned volumes must belong to a clone pool, either one of the preconfigured clone pools or one you create. If you do not select a clone pool, the Solstice Backup server automatically selects the Default Clone Pool.

5. Click Start to begin the cloning.

   Once the cloning process begins, its progress appears in the Status attribute of the Clone Volume Status window.

▼ To View Clone Volume Details

You can view the details of a clone volume, including its estimated capacity, amount of space used, number of files, and mode, by opening the Volume Details window.

To view clone volume details:

1. Select Media>Clone Volume to open the Clone Volume dialog box.

2. Select the volume whose details you want to view in the Volumes attribute.

3. Click Details to open the Volume Details window.

Save Set Cloning

Cloning save sets allows you to create identical copies of backed-up data quickly and easily. The save set cloning feature provides a high degree of flexibility for determining which data to copy. You can clone save sets automatically or manually. You can also clone a single save set, or configure the software to clone only full backups for an entire group.

Overview

The Solstice Backup cloning function ensures the safety of your data by making a copy of each selected save set and writing it to a destination volume. Source and destination media can be of different types and capacities. If necessary, cloned save sets span destination volumes just as they span source volumes. Save sets that fit on one original volume may spill over onto two clone volumes simply because the destination volumes may be a different size or may contain different data than the original. However, the opposite is also possible: the original save sets might require two volumes, while the cloned save sets require only one volume.

Retention Policies

By default, save sets on the source and destination volumes have the same retention policies. However, you may override the preselected retention policy for the cloned volume by changing the mode of the volume to manual recycle in the Volumes window.

When you apply the manual recycle mode to a volume, Solstice Backup software does not overwrite the volume unless you manually change the mode again to recyclable.

Performance Levels

To maintain backup and cloning performance levels, the Solstice Backup server can be configured to clone save sets automatically after completing a scheduled backup. You should allow time to complete the cloning process before scheduling another backup.

Online Index Entries

The Solstice Backup server creates index entries for the cloned save sets and volumes in the online media index. However, both the original and clone volumes use the same set of entries in the online client file index. If you eliminate either the clone or the original volume from the media database, the Solstice Backup server does not let you remove the client file index entries as long as one of the volumes still exists.

Devices

You must have at least two storage devices to clone data. One storage device is used for the source volume and the other for the destination volume.

- If you use libraries with multiple devices, the Solstice Backup server mounts the volumes required for cloning automatically.
- If you use stand-alone devices, you need to mount the volumes manually. The Solstice Backup server displays a message in the Messages window to tell you which volumes to mount. If you want to keep cloned data separate from other cloned data, premount a blank destination volume.

How Save Set Cloning Works

There are two different methods for cloning save set data:

- *Manual* cloning of save sets is performed by entering the appropriate information in the Save Set Clone dialog box. The next section describes this process.
- *Automatic* cloning of save sets can be set in the Group resource. For more information, see "Automatically Cloning a Save Set" on page 350.

▼ To Manually Clone a Save Set

To manually clone a save set, you must first query the Solstice Backup server online databases to locate the correct save set. After querying the database, select the save set and begin the cloning operation.

To manually clone a save set:

1. In the Solstice Backup Administrator program, select Save Set>Clone to open the Save Set Clone dialog box.

Note – The fields in this dialog box, shown in FIGURE 9-2 on page 348, are case sensitive.
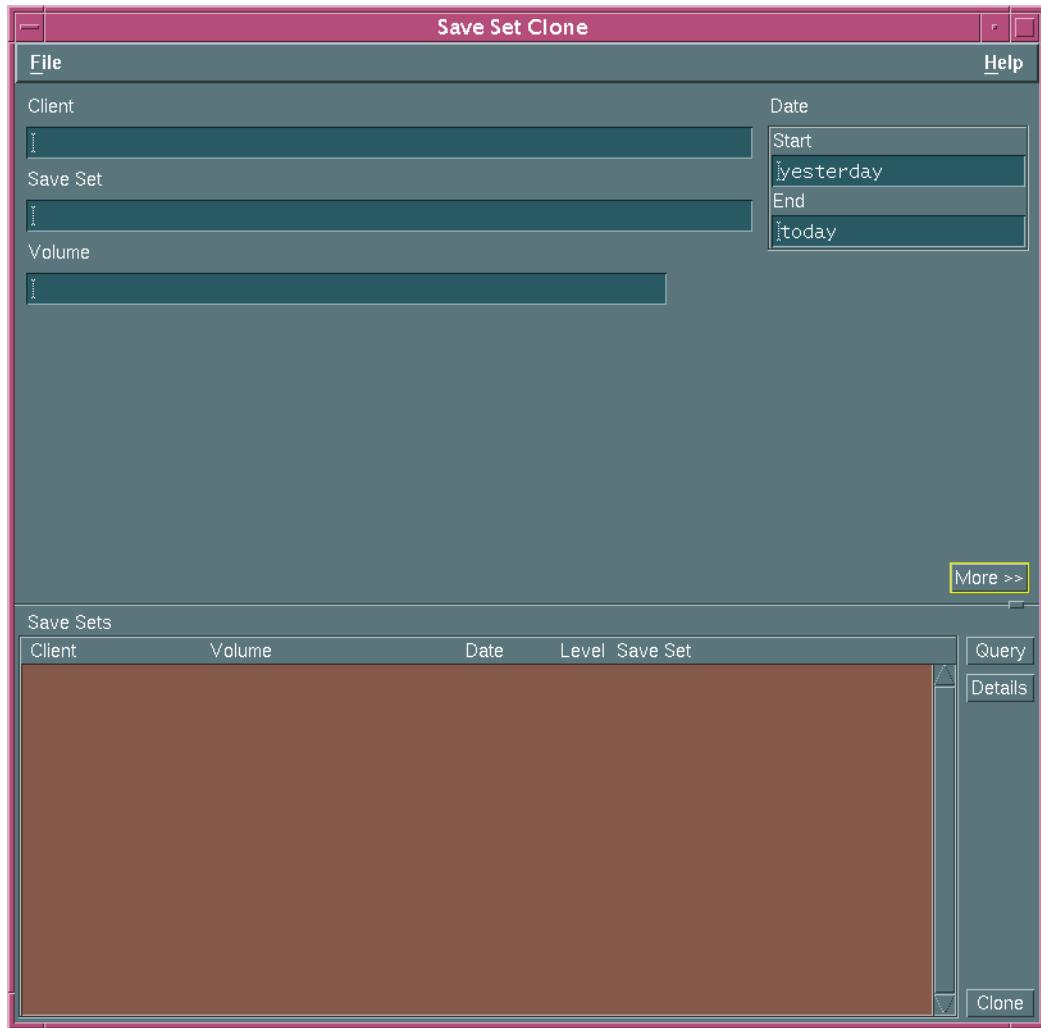
FIGURE 9-2   Save Set Clone Dialog Box

2. For the Client attribute, enter the name of the client associated with the save set.

   If you do not enter a client name, Solstice Backup software locates and displays cloned save sets for all clients.

3. (Optional) For the Save Set attribute, enter the name of the save set.

4. (Optional) For the Volume attribute, enter the name of a volume.

5. (Optional) For the Pool attribute, enter the name of a volume pool.

6. For the Start Date and End Date attributes, enter dates corresponding to when you believe the save set was created. By default, yesterday's date is used for the Start Date and today's date is used for the End Date. You may change these dates by entering dates in any of the following formats:

- Written out completely (for example, November 1, 2001)

- Numerically as mm/dd/yy (for example, 7/1/01)

- As a phrase (for example, two weeks ago, three days ago)

7. Click More to display additional attributes in the Save Set Clone window.

8. (Optional) Select the appropriate Status attributes for the save sets. Browsable, Recyclable, Scanned-in, and Recoverable are already chosen by default. The Solstice Backup software uses seven save set status variations:

TABLE 9-1 Status Attributes

Attribute	Meaning
Browsable	Select if the save set still has an entry in the online file index.
Recyclable	Select if all save sets have passed both the browse and retention policy time periods, and the volume is now available for automatic relabeling and overwriting.
Scanned-in	Select if the save set was scanned in using the scanner command.
Recoverable	Select if the entry for the save set has been removed from the client file index, but is still available for recovery from the media.
In-progress	Select if the save set is currently in the process of being backed up.
Aborted	Select if the save set was either aborted manually by the administrator during a backup, or because the computer crashed.
Suspect	Select if a previous attempt to recover the save set failed.

9. (Optional) Select the appropriate level for the Maximum Level attribute to specify the backup levels of the save sets to be included in the query.

Incremental is specified by default. The Solstice Backup server searches for and displays all the levels up to and including your level choice. For example:

- If you select level *5*, the software searches for and displays save sets backed up at levels Full, 1, 2, 3, 4, and 5.

- If you select level *Full*, the software searches for and displays only save sets backed up at level Full.

- If you select *Incremental*, the software searches for and displays save sets for all levels.

You must always have at least one level selected, and you may not select more than one level at a time.

10. (Optional) Use the Sort by attribute to select what order the save sets appear in the Save Sets list box.

11. Click Query. The save sets that fit your filtering criteria appear in the Save Sets list box.

    You can view details of a selected save set by clicking Details to open the Save Set Clone Details dialog box. This dialog box provides information about the selected save set, such as the name and date, the number of files it contains, and its size and status. Click Close to close the dialog box.

12. Select the save sets you want to clone from the Save Set list.

13. Click Clone. The Save Set Clone Status window appears, showing information about the status of the cloning operation.

Automatically Cloning a Save Set

The Solstice Backup server can automatically clone save set data at the time of a scheduled group backup. This is configured in the Group resource. For general information on configuring groups, including cloning, see "Backup Groups" on page 256.

The backed-up save sets you identify for automatic cloning are cloned without any direct involvement by the administrator. Whenever you know of a save set cloning requirement in advance, you should try to use automatic save set cloning.

Examples of Cloning Save Sets

This section contains three different examples of save set cloning:

- "Example: Manually Cloning a Save Set" on page 351
- "Example: Automatically Cloning One Client" on page 352
- "Example: Automatically Cloning Several Clients" on page 353

Note – All of these examples assume that there are two storage devices containing mounted volumes: one for the source volume and another for the destination (clone) volume.

Example: Manually Cloning a Save Set

In this example, a user has requested that the Solstice Backup administrator manually clone several save sets that are not included in an automatic cloning schedule.

The user needs to send the data to another company located out of state that also uses Solstice Backup software. The Solstice Backup administrator needs to clone the most recent full backup, and any incrementals since the last full backup, to make sure the most current data is sent.

The company receiving the cloned data needs to use the scanner command to retrieve the data. For more information about the scanner command, see "The scanner Program" on page 318.

To clone the save set, the Solstice Backup administrator must have the following information:

- Solstice Backup client name
- Name of the save set
- Date the data was backed up

To manually clone the save sets, the administrator must:

1. In the Solstice Backup Administrator program, select Save Set>Clone to open the Save Set Clone dialog box.

2. For the Client attribute, enter the client name.

3. For the Save Set attribute, enter the save set name.

4. For the Start Date and End Date attributes, enter the dates used for the search.

Note – In this case, there is no need for the administrator to change or select any status choices other than the defaults.

---

5. Click Query.

   All save sets that meet the browsing criteria appear in the Save Sets list box.

6. Select the save sets to clone.

7. Click Details to determine the size of the data and the original volume that contains the data to be cloned.

8. Mount the original volume.

9. Click Clone.

Example: Automatically Cloning One Client

In this example, a system administrator wants to clone data automatically, on a daily basis, for the Solstice Backup client computer jupiter, where the company product documentation resides. Client jupiter performs an incremental backup every night of the week except Friday, when it performs a full backup.

The system administrator needs to perform the following tasks:

- Create a group named Documents to which the client jupiter will belong.
- Make the appropriate selections for the client jupiter, which includes selecting the group Documents.

To automatically clone jupiter's data, the system administrator must:

1. Create the group Documents.

   For more information on creating groups, see "To Create a Group" on page 260.

2. Select the appropriate choices in the Groups dialog box. The required choices are:
   - Autostart = Enabled
   - Clones = Yes
   - Clone pool = Default Clone (or a custom clone pool)

3. Create the client jupiter.

   For more information on creating clients, see "To Create a Client" on page 141.

4. Select the appropriate backup choices for jupiter. The required choices are:
   - Schedule = Full Every Friday
   - Group = Documents

With this set of instructions, the administrator creates a group for just one client, so that only that client's data is cloned on a daily basis.

Example: Automatically Cloning Several Clients

illustrates the level of control you can achieve when cloning data with Solstice Backup software. You can create an unlimited number of automatic cloning configurations using the Group, Schedule, and Client resources.
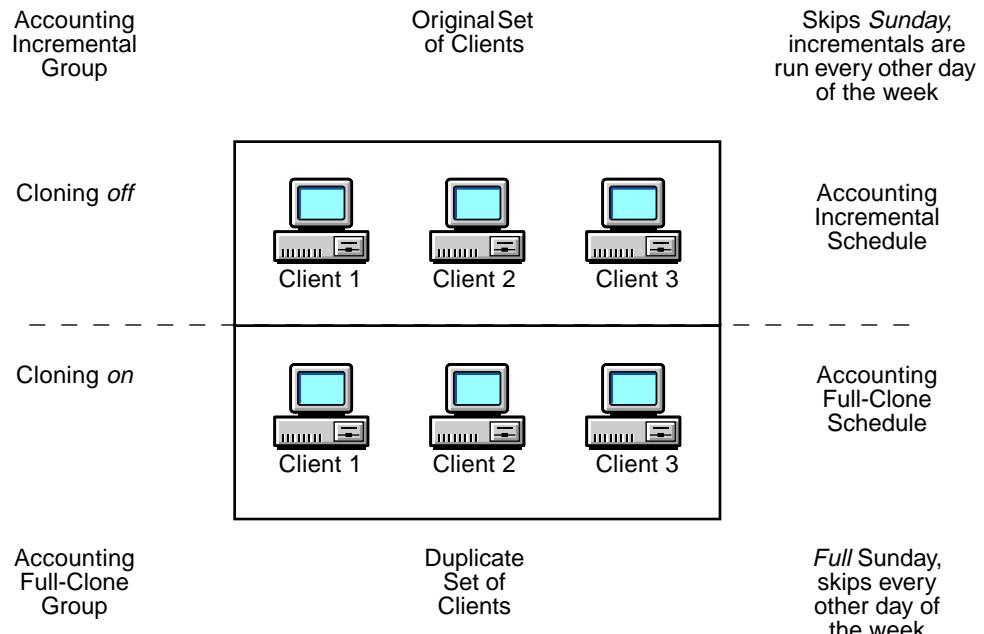
FIGURE 9-3   Automatic Cloning Scheme

In this example, the system administrator wants to automatically clone the data for a small group of Solstice Backup clients with valuable accounting information. However, only clones of full backups are needed, to store in an offsite vault.

To do this, two different groups and schedules are created and applied to two identical sets of individual clients. Both groups belong to the same pool and use the browse and retention policies specified for the Client resource.

Note – To keep full backups and incremental backups on separate volumes, the administrator would also need to create two pools.

The system administrator needs to accomplish the following tasks:

■ Create two schedules, one to perform incremental backups for the clients and one to perform full backups and clones:

– Create one schedule named Accounting Incremental (shown in TABLE 9-2).

TABLE 9-2   Accounting Incremental Schedule

Sun	Mon	Tue	Wed	Thu	Fri	Sat
skip	incr	incr	incr	incr	incr	incr

– Create another schedule named Accounting Full-Clone (shown in TABLE 9-3).

TABLE 9-3   Accounting Full-Clone Schedule

Sun	Mon	Tue	Wed	Thu	Fri	Sat
full	skip	skip	skip	skip	skip	skip

Note – By creating two schedules and selecting them for the appropriate set of clients, the system administrator avoids cloning all of the incremental backups.

■ Create two groups, one named Accounting Incremental and another named Accounting Full-Clone.

Note – Automatic cloning is applied to one of the groups. As a result, two groups are necessary in this situation. The Accounting Incremental group is used for performing incremental backups, and the Accounting Full-Clone group is used for performing full backups and clones.

■ Create *two* clients with the same name for each client computer. Two clients of the same name are needed because only one schedule can be applied to a Solstice Backup client at a time. In this example, one set of clients uses the Accounting Incremental schedule and the other (identical) set of clients uses the Accounting Full-Clone schedule.

■ Select the appropriate group and schedule for each set of clients.

To automatically clone data in this example, the system administrator must:

1. In the Solstice Backup Administrator program, select Configure>Schedules to open the Schedules window.

2. Create the Accounting Incremental schedule to perform incremental backups every day except Sunday, when it skips the backup (see TABLE 9-2 on page 354).

   For more details on creating schedules, see "To Create a Schedule" on page 276.

3. Create the Accounting Full-Clone schedule to perform full backups every Sunday and skip backups the rest of the week (see TABLE 9-3 on page 354).

4. Create the Accounting Incremental group. The required attributes are:

- Autostart = Enabled
- Clones = No

For instructions on creating groups, see "To Create a Group" on page 260.

5. Create the Accounting Full-Clone group. The required attributes are:

- Autostart = Enabled
- Clones = Yes
- Clone = Backup Clone

6. Create two sets of identical clients, and select the appropriate schedule and group for each client.

   For more information on creating clients, see "To Create a Client" on page 141.

7. Select the first set of clients, one client at a time, and assign them to the Accounting Incremental group. Make the appropriate selections in the Create Client dialog box. The required attribute are:

- Schedule = Accounting Incremental
- Group = Accounting Incremental

   After making these choices *for each client*, click OK.

8. Select the duplicate set of clients, one client at a time, this time selecting Accounting Full-Clone for the Schedule and Group attributes.

   After making these choices *for each client*, click OK.

When the Accounting Full-Clone group runs its full backup on Sunday, the data is automatically cloned after the backup is completed. The incremental backups (skipped every other day of the week) are *not* cloned.

Recovering Cloned Save Sets

When recovering cloned save sets, Solstice Backup software always looks for the original volumes by default. The volumes required for a save set recovery appear in the Required Volumes dialog box.

In order to recover save set files from a clone volume, rather than the original volume, you must first designate the original save set as *suspect*. Once the Solstice Backup server determines that the original save set is marked as suspect, it looks for a cloned version. If a clone volume containing the cloned save set exists, it appears in the Required Volumes dialog box. Otherwise, the software attempts to recover the data from the suspect volume.

If you cannot determine which specific save set on the volume is the one you want to recover (in order to make it suspect), you should change the status of *all* of the save sets on the volume to suspect. This ensures that Solstice Backup software cannot find the save set it needs on the original volume during the recovery process.

Note – The Solstice Backup server does not display more than one set of volumes needed for recovery at one time. Therefore, although there may be more than one set of clones available to you for recovery, only one set at a time appears in the Required Volumes dialog box.

For information about changing the status of a save set, see "To Change the Status of a Save Set" on page 155.

Other Information and Procedures

This section contains miscellaneous information and procedures that may be useful when you are cloning data with Solstice Backup software.

▼ To Clone Archived Data

To clone an existing archive volume:

1. In the Solstice Backup Administrator program, select Save Set>Clone to open the Save Set Clone dialog box.

2. In the Save Set Clone dialog box, make your selections for browsing save sets.

   For more information, see "To Manually Clone a Save Set" on page 347.

3. In the Pool text box, enter Archive.

4. Click Query to display the save sets matching your criteria.

5. Select the save sets you wish to clone.

6. Click Clone.

Clone Storage Node Affinity

The link between a client's resource of a storage node and a list of available storage nodes to receive cloned save sets from the storage node client is called *clone storage node affinity*. Data is cloned from media that contains the original save sets to media on the specified clone storage node. You define clone storage node affinity in the Clone Storage Nodes attribute, which is found in the Client resource of a storage node. When you make a change, the changed value is propagated to any additional Client resources configured for that storage node client.

The Clone Storage Nodes attribute allows you to specify a network interface for storage nodes that perform cloning operations different than the one specified for the storage node's remote device. The server uses the exact hostname you specify for the attribute, instead of using the hostname prefix for the remote device name configured in the Device resource.

When a volume is being cloned, the Solstice Backup server checks the value of the Clone Storage Nodes attribute for that storage node client:

■ If the Clone Storage Nodes attribute is a null value, then the value listed for the server's Clone Storage Nodes attribute is used.

■ If that list also has a null value, then the server's Storage Node attribute is used.

Compatibility is maintained with the existing clone function that follows the server's Storage Node attribute.

To independently direct clones from each storage node, add the hostname of the storage node that you want to receive the directed clones to the Clone Storage Nodes attribute in the Client resource configured for the storage node. The first entry made on the list that has a functional, enabled device is selected to receive the cloned data from the storage node.

To direct clones from all storage nodes to the same destination, leave the Clone Storage Nodes attribute blank for the Client resources you configure for the storage nodes, and configure only the Solstice Backup server's Clone Storage Nodes attribute. This provides a single source of control for clone destination.

The client file index and media database entries for the save sets cloned to media on a remote device on a storage node still reside on the Solstice Backup server. This enforces the browse and retention policies in the same manner as for any cloned save sets that reside on the media in a device that is locally attached to the server.

Autochangers and Silos

This chapter provides information about how to configure and operate the Solstice Backup server with autochangers and silos. Working with autochangers and silos requires knowledge of installing and setting up storage devices. For additional information pertinent to all devices, including mounting and unmounting volumes, managing media, using pools, and cloning volumes, see Chapter 5.

Using Solstice Backup software with an autochanger requires you to purchase and enable the Solstice Backup Autochanger Software Module. Contact your Solstice Backup sales representative for more information. For general information on licensing Solstice Backup products, see the Solstice Backup Roadmap.

This main sections in this chapter are:

- "Autochanger Installation and Configuration" on page 360
- "Dynamic Drive Sharing" on page 374
- "Sharing Autochangers between Solstice Backup Hosts" on page 379
- "Autochanger Operations" on page 387
- "Labeling Volumes with Bar Codes" on page 396
- "Using Autochangers with a Volume Import/Export Capability" on page 401
- "Tips and Suggestions for Using Autochangers" on page 403
- "Silos" on page 407
- "Media Management in a Silo" on page 413

Autochanger Installation and Configuration

Autochangers automate the task of loading, mounting, and labeling backup media. The term *autochanger* refers to a variety of storage devices, including the following: jukebox, carousel, library, near-line storage, datawheel, and autoloader. Note that tape "stackers" are not considered autochangers by the Solstice Backup software. Stackers access tapes sequentially, while autochangers can access media in any order.

Before you back up to an autochanger, perform the following tasks:

1. Install and enable the Solstice Backup software.

2. Install and enable the Autochanger Software Module.

3. Load and label your volumes.

Some autochanger configuration occurs when you install and enable the Autochanger Software Module. After installing the autochanger, you can make changes to the available slots, devices, bar code choices, auto media management, and cleaning devices.

The Solstice Backup software does not support emulation mode for any storage device. It communicates with devices only in their native mode. For example, do not run an ADIC DLT in Exabyte 440 mode.

This section includes the following information:

▼ To Add Autochangers

You cannot add or create autochangers using the Solstice Backup Administrator program; instead you must use the jbconfig program at the command line. However, once the autochanger has been added and configured, you can use the Solstice Backup Administrator program to modify its configuration.

Caution – Prior to running the jbconfig program, determine what existing devices are defined in the Devices panel of the Solstice Backup Administrator program. If it does not show the correct media type, then delete this device. For information on deleting devices, see "To Delete a Device" on page 222.

If you do not perform this task, the Solstice Backup server backs up to the device using the wrong blocking factor. While adding the autochanger using the jbconfig program, you will be able to specify the correct devices your autochanger contains.

To add autochangers:

1. Navigate to the directory containing your Solstice Backup software.

2. Enter the following command at the prompt:

```
# jbconfig
```

3. The jbconfig program displays the following installation script and asks for your choices. Press [Return] to accept the suggestions.

Example: Configuring a SCSI Autochanger on a Solaris System

The following example shows the required responses to configure a SCSI autochanger on a Solaris system. The driver software detects and displays the information for all SCSI autochangers attached to the system.

```
# jbconfig

1) Install a SmartMedia Jukebox.

2) Install an Autodetected SCSI Jukebox.

3) Install an Autodetected NDMP SCSI Jukebox.

4) Install an SJI Jukebox.

5) Install an STL Silo.

What kind of Jukebox are you installing? [1] 2

These are the SCSI Jukeboxes currently attached to your system:

1) scsidev@1.2.0: DLI Libra Series
2) scsidev@0.2.1: Quantum DLT/Compaq DLT

Which one do you want to install? 2

Installing a 'Quantum DLT/Compaq DLT' jukebox.

Name you would like to assign to the jukebox device? QuantumDLT_1

Pathname of the control port for the jukebox device?
[scsidev@0.2.1] [Return]

Do you want automated device cleaning support enabled? (yes/no) n

Enter pathname of media drive 1 [/dev/nrst8]:? [Return]

This media device has not been configured yet. Please select a
media device type for /dev/nrst8.
```

```
a) himt
b) qic
c) 4mm
d) 8mm
e) 8mm 5GB
f) 3480
g) dlt
h) vhs
i) optical

Choice? g


Jukebox has been added successfully
```

Example: Configuring an SJI Autochanger

The following example shows the responses to configure an SJI autochanger.

```
# jbconfig

1) Install a SmartMedia Jukebox.

2) Install an Autodetected SCSI Jukebox.

3) Install an Autodetected NDMP SCSI Jukebox.

4) Install an SJI Jukebox.

5) Install an STL Silo.

What kind of Jukebox are you installing? [1] 3

Enter the number corresponding to the type of jukebox you are
installing:

1) ADIC-1200c/ADIC-1200d          22) Qualstar
2) ADIC-VLS                       23) Spectralogic
3) ARC DiamondBack                24) STK 9704/Lago 340
4) Sun 20Gb 4mm Tape              25) STK 9708/Lago 380 (SCSI)
5) Breece Hill                    26) StorageTek 9730
6) Breece Hill Saguaro            27) StorageTek 9738
7) Philips Blackjack              28) Dell PowerVault 130T
8) DLI Libra Series               29) Hewlett-Packard A4853A
9) Quantum DLT/Compaq DLT         30) IBM 3570
.
.
.

Choice? 9

Installing a 'Quantum DLT/Compaq DLT' jukebox.

Name you would like to assign to the jukebox device? dlt

Pathname of the control port for the jukebox device? scsidev@1.2.0
[Return]

Do you want automated device cleaning support enabled? (yes/no) n
```

```
Enter pathname of media drive 1 [/dev/nrst8]:? [Return]

This media device has not been configured yet. Please select a
media device type for /dev/nrst8.

a) himt
b) qic
c) 4mm
d) 8mm
e) 8mm 5GB
f) 3480
g) dlt
h) vhs
i) optical

Choice? c

Jukebox has been added successfully.
```

▼ To Test the Autochanger Connection

After configuring the autochanger, you should test the connection. To test the
autochanger connection:

1. Log in as root on the Solstice Backup server or storage node.

2. Insert two volumes, one each into the first and last slots of the autochanger. Make
   sure that the devices are empty and that any device doors are open.

3. Enter the following command at the prompt:

```
# jbexercise
```

You need to specify the control port and the device type. The control port for SCSI autochanger models is typically expressed in the format scsidev@n.n.n. You can obtain the exact control port pathname from the response displayed by the jbconfig command script:

```
These are the SCSI Jukeboxes currently attached to your system:

1) scsidev@1.2.0: DLI Libra Series

2) scsidev@0.2.1: Quantum DLT/Compaq DLT
```

For example, the following command runs the jbexercise program on the Quantum DLT/Compaq DLT autochanger detected by the jbconfig command:

```
# jbexercise -c scsidev@0.2.1 -m "Quantum DLT/Compaq DLT"
```

For more information about jbexercise, refer to the *Legato Command Reference Guide* or to the jbexercise man page.

Testing Device Configuration Using the jbverify Program

You can check the devices that are defined in the Solstice Backup database using the jbverify program. This program ensures that each device is configured properly by checking for accessibility and usability.

By default, jbverify **checks all the devices in the Solstice Backup database. However, the following command-line options can be used to limit the devices checked by the command:**

- Autochangers — using the -j option
- Stand-alone drives that are not contained in autochangers — using the -d option
- Local devices — using the -l option
- Specific jukeboxes — using the -J option
- Specific drives — using the -D option

Note – The jbverify program does not test Legato GEMS SmartMedia devices.

For further information regarding the jbverify program, refer to the *Legato Command Reference Guide* or to the jbverify man page.

Requirements for Running the jbverify Program

The following sections describe the requirements of the jbverify program:

Requirements for Servers and Clients

For the jbverify program to work, the Solstice Backup server process (nsrd) must be running on the server computer and the client processes (nsrexecd) must be running on the client computers.

When the jbverify program is run on a server, it can test any device that has been configured in the Solstice Backup software.

Requirements for Storage Nodes

The jbverify program can be run on any storage node and can be used to test any device on that storage node. The device, however, must have been configured in the Solstice Backup software.

The jbverify program can be used to remotely test devices from a storage node.

In such a case, you must complete one of the following procedures:

- The target storage node must have the invoking storage node listed in its Servers file.
- The nsrexecd program on the target storage node must be started with the -s option. The invoking storage node must be the argument. For example:

```
# nsrexecd -s invoking_storage_node -s Solstice Backup_Server
```

▼ To Test Autochangers

For autochangers, the jbverify program performs the load and unload operations on each drive without performing any write operations, unless the -t option is used. If no slot is specified on the command line, the jbverify program checks all the slots that are available to the Solstice Backup server and then loads the first available slot.

To test only autochangers, enter the jbverify command with the -j option:

```
# jbverify -j jukebox_name
```

To use a specific slot for the test, use the -S option:

```
# jbverify -S slot
```

To test a specific jukebox, use the -J option:

```
# jbverify -J jukebox_name
```

▼ To Test Stand-Alone Devices

To test stand-alone drives, without testing any other devices, use the -d option:

```
# jbverify -d device_name
```

To test a specific drive, use the -D option:

```
# jbverify -D device_name
```

▼ To Test Local Devices

To test local devices, without testing any other devices, use the -l option.

```
# jbverify -l
```

▼ To Test Tapes with the jbverify Program

The jbverify program can test tapes that are loaded into the drives that are being tested. The program tests these tapes by using the Solstice Backup tape exercise program (tapeexercise **on UNIX or** tapeexer **on Microsoft Windows).**

For further information regarding the tapeexercise and tapeexer programs, refer to the *Legato Command Reference Guide* or to the tapeexercise man page (UNIX platforms only).

To test a tape using the jbverify program, use the -t option.

Configuring Autochanger Attributes

Attributes for a specific autochanger are listed in the Jukeboxes resource. Select Media>Jukeboxes to display the Jukeboxes dialog box. Autochangers configured using jbconfig are listed in the Jukeboxes attribute in this dialog box.

Key Autochanger Attributes

TABLE 10-1 lists some of the key autochanger attributes and their functions.

TABLE 10-1   Autochanger Attributes

Attribute	Action
Control Port	Enter the pathname of the autochanger's command control port.
Devices	Enter the names of each device in the autochanger.
Bar Code Reader	Select Yes to enable this function for an autochanger with bar code reading capability.
Match Bar Code Labels	Select Yes for your bar code-enabled autochanger to match bar code labels and volume labels.
Cleaning Slots	Enter the slot number containing the cleaning cartridge.
Default Cleanings	Enter the number of cleanings the cleaning cartridge can handle before the Solstice Backup server sends you a notification to replace the cartridge.
Auto Clean	Select Yes to enable autocleaning.
Auto Media Management	Select Yes to enable auto media management, which includes automatic labeling of blank volumes.
Available Slots	Enter the range of autochanger slots available for the Solstice Backup server to use.

▼ To Add Devices to an Autochanger

To add more devices to an autochanger, you must first configure the Solstice Backup server to recognize the added devices. The Device resource contains attributes for you to supply information about the additional autochanger device. Then you must enter the pathname for the autochanger device in the Jukebox resource so the server can recognize that the device belongs to the autochanger. For an example of pathname naming, see FIGURE 10-1 on page 370.

To add devices to an autochanger:

1. In the Solstice Backup Administrator, select Devices from the Media menu to open the Devices window.

2. Click Create.

3. Enter the autochanger device pathname in the Name attribute.

   If an autochanger has more than one device, you must list the device pathnames in the same order as their physical location in the autochanger. If you are unsure of their order, do one of the following:

   ■ Refer to the autochanger hardware manual.

   ■ Log in as root on the Solstice Backup server or storage node computer and enter the `inquire` command at the shell prompt. The `inquire` command returns a list of the SCSI devices attached to your system.

4. Select the appropriate media type in the Media Type attribute.

5. Complete the other attributes as necessary and click Apply. For more details on configuring devices, see "To Add a Device" on page 220.

6. Select Media>Jukeboxes to open the Jukeboxes resource.

7. Select the appropriate jukebox in the Jukeboxes attribute.

8. Enter the storage device pathname in the Devices attribute and click Add.

If you like, you can make other modifications to your autochanger configurations at this time. For example, you might select bar code labels, or decide to change the range of available slots.

FIGURE 10-1 on page 370 shows an example of pathname naming.
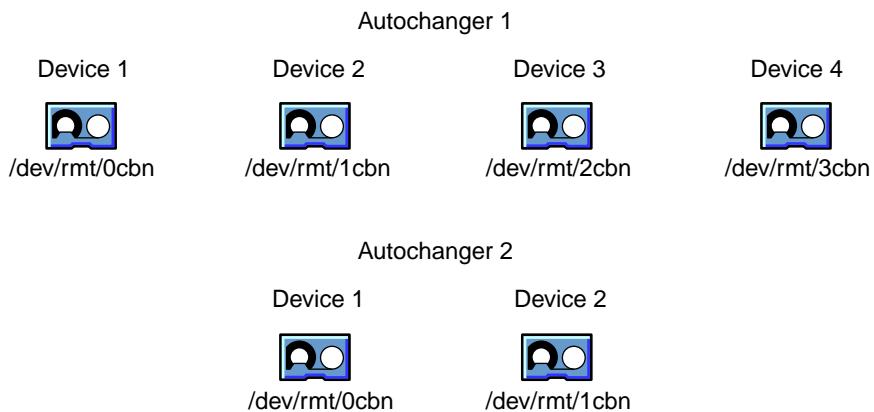
FIGURE 10-1 Autochanger Device Pathnames

Caution – You need to reset an autochanger every time you change or modify the device. This should be done with the `sjiielm` utility and/or preferably with the `nsrjb -HE` command-line program. For example, the autochanger inventory must be correct after adding drives to an SJI-compliant autochanger, such as adding DLT7000 drives to an ETL 7/3500 device.

To make Solstice Backup software aware of these new drives, execute `nsrjb -HE` to reset the autochanger. This must be done after running `jbconfig`. The `-E` option reinitializes the autochanger's element status. Some autochangers can keep track of whether there is media in a component in the autochanger. This feature is known as an *element status* capability.

In some cases, there is no need for a manual hardware reset if the autochanger is able to reinventory itself.

▼ To Edit an Autochanger

To edit an autochanger:

1. Select Media>Jukeboxes to open the Jukeboxes window.

2. Select the autochanger from the list displayed in the Jukeboxes attribute.

3. Edit the attributes of the autochanger and click Apply.

▼ To Delete an Autochanger

To delete an autochanger:

1. Select Media>Jukeboxes to open the Jukeboxes window.

2. Select the autochanger from the list displayed in the Jukeboxes attribute.

3. Click Delete.

   When prompted to confirm that you really want to delete that autochanger, Click OK.

4. Click Apply.

If you want to reinstall the autochanger or add a new autochanger to your Solstice Backup server, you need to run the `jbconfig` program.

Remote Autochanger Management

Autochangers that are connected to storage nodes require a few additional configuration and management steps.

You can control most operations on remote autochanger devices from the Solstice Backup Administration program. But for some remote autochanger operations, such as reset, you must use the nsrjb or jbconfig commands on the storage node. You can issue the commands as root, either from the local computer or through a remote login session.

After you install the storage node binaries on the storage node, define the storage node's devices. The method for defining devices is described in "To Configure a Storage Node" on page 237. An overview of how to add autochanger devices is provided here.

▼ To to Add a Remove Autochanger Device

To add a remote autochanger device:

1. In the Solstice Backup Administrator program, select Setup Server from the Server menu.

2. Add the storage node's hostname to the Administrator attribute in the Server resource by using the following syntax:

```
root@storage_node_hostname
```

3. Run the jbconfig program on the storage node computer to define each device in the autochanger.

For more information, see "Autochanger Installation and Configuration" on page 360, and refer to the *Legato Command Reference Guide* or the jbconfig man page.

The device names for remote devices begin with a prefix of "rd=" and the name of the storage node computer. For example, rd=jupiter:/dev/rmt/1mbn is a device called /dev/rmt/1mbn on a storage node computer called jupiter.

▼ To Adjust Sleep Periods for Autochanger Tasks

The Jukebox resource includes attributes that specify the number of seconds that the autochanger is inactive after volume loading, unloading, or ejecting. This function is known as *sleeping*, and during a sleep period, the autochanger is restricted from receiving or performing other operations.

This period of inactivity is sometimes necessary when loading, unloading, and ejecting volumes on some autochangers whose devices take longer than the usual time to perform the task. Without the sleep period in effect, the loading or unloading of volumes might fail.

When you configure a Jukebox resource using the `jbconfig` program, Solstice Backup software automatically configures default sleeping periods. You should only change these values if you are troubleshooting your autochanger's performance, or if a Solstice Backup technical support specialist instructs you. Typically, the higher the sleep values specified in the attributes, the longer it will take the autochanger to perform the task. Be cautious when changing these values.

TABLE 10-2 lists the attributes related to the autochanger sleep function.

TABLE 10-2   Jukebox Resource Sleep Attributes

Attribute	Description
Eject Sleep	Specifies the number of seconds the Solstice Backup software sleeps after ejecting a volume. The Solstice Backup software uses this attribute if the autochanger is not configured for volume auto-ejection. Autochangers configured for volume auto-ejection do not perform a separate eject operation, so this attribute has no effect on them.
Unload Sleep	Specifies the number of seconds an autochanger sleeps between unloading a volume and moving the volume from the device to a slot. The Solstice Backup software uses this attribute for all autochangers when a volume is unloaded.
Load Sleep	Specifies the number of seconds the Solstice Backup software sleeps between loading a volume and trying to access that volume's data.

To access and set the sleep-related Jukebox resource attributes:

1. Select Media>Jukeboxes to open the Jukeboxes window.

2. Select View>Details to display hidden attributes.

3. Adjust one or more of the sleep-related attributes.

Dynamic Drive Sharing

Dynamic Drive Sharing (DDS) is a feature that provides Solstice Backup software with the ability to recognize shared drives. DDS enables Solstice Backup software to skip the shared drives that are in use and route the backups or recoveries to other available shared drives.

This section contains the following topics:

- "To Enable DDS When Configuring an Autochanger" on page 374
- "To Stop Sharing a Device within an Autochanger" on page 378
- "To Change the Idle Device Timeout Attribute" on page 379

For further information about the DDS feature, see "Dynamic Drive Sharing" on page 225.

For information explaining how to enable DDS with the jbconfig program, see "To Enable DDS When Configuring an Autochanger" on page 374.

Changes to the jbconfig Program

Changes have been made to the jbconfig program that allow you to enable DDS when you are configuring an autochanger. The jbconfig program now prompts you to do:

- Select a drive to share.
- Assign the hardware ID of the devices that will be sharing this drive.

▼ To Enable DDS When Configuring an Autochanger

You can enable the DDS feature when you are configuring an autochanger through the jbconfig program. A number of changes have been made to the program, enabling you to configure the resources to be used with the DDS feature.

Note – Before you enable DDS through the jbconfig program, you must ensure that each node (or server) that is to be sharing a drive is licensed. Without licensing before configuration, the jbconfig program fails to configure the autochanger. For information explaining how to license a product, the Solstice Backup Roadmap.

To enable DDS when you configure an autochanger:

1. Enter the following command at the prompt:

```
# jbconfig
```

2. Select the autochanger appropriate for your system and complete the prompts as required.

3. Enter yes when the script prompts whether the drive will be shared.

4. Assign a hardware ID when the script prompts you for one.

   For further information about the hardware ID, see "Hardware ID Attribute" on page 229.

5. Complete the prompts as required for the device that will be accessing the shared drive.

6. Enter the name of the second device that will be sharing the drive when prompted.

7. Complete the prompts for the second device.

Sample jbconfig Output

This sample jbconfig output is for the following configuration:

■ Server (jupiter)

■ A storage node (sn_1)

■ A shared IBM 3590 jukebox with one drive (drive_1)

```
# jbconfig

        1) Install a SmartMedia Jukebox.
        2) Install an Autodetected SCSI Jukebox.
        3) Install an Autodetected NDMP SCSI Jukebox.
        4) Install an SJI Jukebox.
        5) Install an STL Silo.

What kind of Jukebox are you installing? [1] 2

These are the SCSI Jukeboxes currently attached to your system:
1) scsidev@2.3.1: Standard SCSI Jukebox, Vendor <IBM>, Product
<03590E11>
```

```
Which one do you want to install? 1

Installing an 'Standard SCSI Jukebox' jukebox.

Name you would like to assign to the jukebox device? IBM_3590

A brief description of the jukebox device? [scsidev@2.3.1]

Pathname of the control port for the jukebox device?
[scsidev@2.3.1]

Do you want automated device cleaning support enabled? (yes/no) n

Will media drive 1 be shared by multiple device paths? (yes/no) y
```

Note – From this point forward, the jbconfig program has been modified to
incorporate changes required to configure DDS.

```
Since this drive will be shared, you will be asked for a hardware
id, which uniquely identifies the drive.  Each device path that
shares the drive will be given the same hardware id value.  When
a device has a hardware id value, its messages will contain a
number that uniquely represents the hardware id.  This number is
visible in administrator programs, such as nwadmin.

Enter hardware id for media drive 1: ? drive_1

Continue with this hardware id? (yes/no) y

This hardware id will only be used when creating new device paths.

Configuring device path 1 for media drive 1 with hw id ` drive_1'

Enter hostname that owns device path 1 for media drive 1: ? jupiter

Enter pathname of device 1 for media drive 1: ? /dev/rmt/0cbn

using '/dev/rmt/0cbn' as device name

Should the drive be configured as a NDMP device? (yes/no) n
```

```
This media device has not been configured yet.  Please
select a media device type for /dev/rmt/0cbn.

        1) 3480                   22) dst
        2) 3570                   23) dst (NT)
        3) 3590                   24) dtf
         .                          .
         .                          .
         .                          .

Choice? 3

Any more device paths for media drive 1? (yes/no) y

Configuring device path 2 for media drive 1 with hw id `drive_1'
```

Note – At the following prompt, enter the name of another host that will also use
the same hardware ID to access the shared drive.

```
Enter hostname that owns device path 2 for media drive 1: ?
[infinity] sn_1

Enter pathname of device 2 for media drive 1: ? /dev/rmt/0cbn

using 'rd=sn_1:/dev/rmt/0cbn' as device name

Should the drive be configured as a NDMP device? (yes/no) n

This media device has not been configured yet.  Please
select a media device type for rd=sn_1:/dev/rmt/0cbn.

        1) 3480                   22) dst
        2) 3570                   23) dst (NT)
        3) 3590                   24) dtf
         .                          .
         .                          .
         .                          .

Choice? 3

Any more device paths for media drive 1? (yes/no) n
```

```
Your jukebox does not report that it has a barcode reader, but it
does report that it can handle volume tags.  Some jukeboxes that
have barcode readers report this way.  If your jukebox has a
barcode reader, do you want barcode reader support enabled?
(yes/no) n

Jukebox has been added successfully

Would you like to configure another jukebox? (yes/no) n
```

▼ To Stop Sharing a Device within an Autochanger

To stop sharing a DDS-configured device within an autochanger:

1. Delete the autochanger:

   a. In the Solstice Backup Administrator program, select Media>Jukeboxes to open
   the Jukeboxes window.

   b. Select the autochanger from the list displayed in the Jukeboxes attribute.

   c. Click Delete.

   d. When prompted to confirm that you really want to delete the autochanger,
   click OK.

   e. Click Apply.

2. Remove the value from the hardware ID attribute in the Device resource.

3. Reconfigure the autochanger using the `jbconfig` program.

Caution – When you stop sharing a DDS-configured drive, this does not mean that
the drive is disabled. This means that Solstice Backup software is no longer aware
that a specified drive is shared.

Consequently, so as not to impede Solstice Backup operations when you stop
sharing a device, you should remove the other device instances.

▼ To Change the Idle Device Timeout Attribute

You might notice that a tape remains mounted in its drive for 10 to 15 minutes after a backup has completed. Other requests for the drive from another device path will wait during this period.

You can modify this timeout value by adjusting the Idle Device Timeout attribute. To access this attribute:

1. In the Solstice Backup Administrator program, select Media>Jukeboxes to open the Jukeboxes window.

2. Select View>Details.

3. Scroll down to Idle Device Timeout.

4. Select a time that is appropriate for your system configuration.

   The default Idle Device Timeout value is 10 minutes.

Note – Selecting zero means that the device never times out. Consequently, you would have to eject the tape manually.

Sharing Autochangers between Solstice Backup Hosts

The Solstice Backup software permits the control and usage of individual devices within an autochanger to different Solstice Backup hosts (a Solstice Backup server or storage node) within a datazone, a feature known as *autochanger sharing*.

The presence of a storage area network (SAN) within your datazone is not required for using autochanger sharing.

Autochanger Sharing Licensing and Availability

Licensing and availability of the Solstice Backup autochanger sharing feature varies with the type of Solstice Backup license you have.

- Solstice Backup Power Edition — Autochanger sharing is automatically enabled as part of Solstice Backup Power Edition. However, you must still have a valid Solstice Backup autochanger license for each autochanger.

- Solstice Backup Network Edition — An enabler license for Solstice Backup autochanger sharing is required to use the autochanger sharing feature within a single Network Edition datazone. Only a single autochanger sharing enabler is required no matter how many autochangers you have within the datazone. However, you must still have a valid Solstice Backup autochanger license for each autochanger.

- Solstice Backup Workgroup Edition — Autochanger sharing is not available for Solstice Backup Workgroup Edition. If your datazone environment requires autochanger sharing, upgrade your Solstice Backup license to Network Edition or Power Edition.

Licensing autochanger sharing is similar to licensing other Legato NetWorker Modules and add-on products. Only a single autochanger sharing license is required for all autochangers in your Solstice Backup datazone, regardless of the number of autochangers present. However, you have to license each autochanger to enable the autochanger functionality itself. For general information on licensing of Solstice Backup products, the Solstice Backup Roadmap.

Note – Only Solstice Backup Network Edition requires you to license autochanger sharing, since this feature is already enabled in Power Edition, and is not available with Workgroup Edition.

How Autochanger Sharing Works

Autochanger sharing allows one Solstice Backup host to control the autochanger's robotic arm, while other Solstice Backup hosts (as well as the host controlling the robotic arm) can each control and use specific autochanger devices. A specific device can be controlled only by a single Solstice Backup host.

FIGURE 10-2 shows how multiple Solstice Backup hosts can share autochanger devices.

Solstice Backup Server

Storage Node

Device 1

Device 2

Shared Autochanger

Legend	
Control of Device	Control of Autochanger
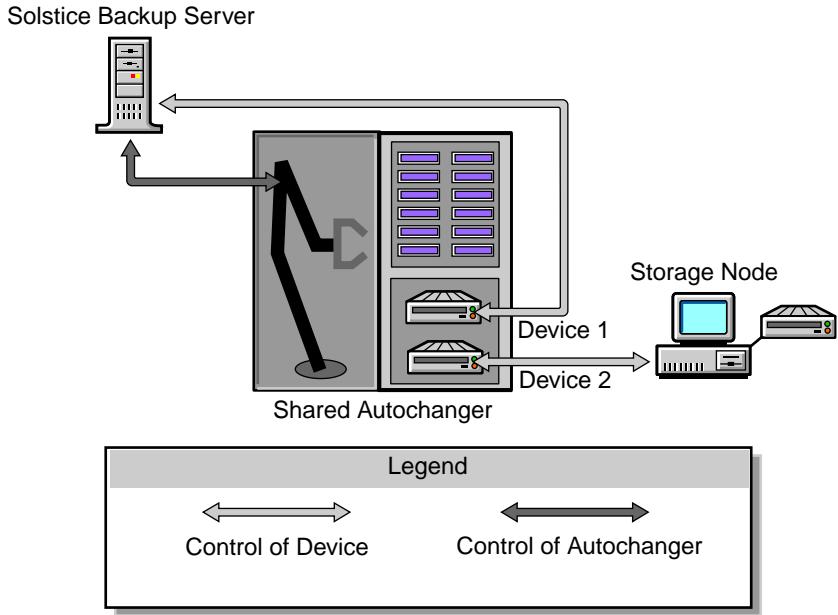
FIGURE 10-2 How Autochanger Sharing Works

Autochanger Sharing Configuration

You can set an autochanger for sharing during the initial autochanger setup using the jbconfig command. The jbconfig command must be run from the Solstice Backup host that controls the autochanger (controls the robotic arm). When running jbconfig, you will be prompted to enter the Solstice Backup hostname controlling the first autochanger device. The default hostname value is the local Solstice Backup host on which jbconfig is running. However, you can change the default hostname value to another Solstice Backup host within the datazone. You will continue to be prompted to enter the controlling host for each device in the autochanger.

Example: The `jbconfig` Command Additions

The following example shows the additions to the `jbconfig` command for the Solstice Backup server `neptune` and a storage node `pluto`. In this example, `jbconfig` is being run on the server `neptune`, which also controls the robotic arm. However, the first media device of the autochanger is being assigned to the storage node `pluto`. The user responses are in bold.

```
Enter hostname that owns media drive 1: ? [neptune] pluto

Enter pathname of media drive 1: ?

/dev/rmt/1cbn

using 'rd=pluto:/dev/rmt/1cbn' as device name
```

Device Naming during Configuration

The `jbconfig` command automatically generates the device names for a shared autochanger according to the following conditions:

- If a device is controlled by the Solstice Backup server, only local device syntax is used (`device_path`).
- If the device is controlled by any other host, the hostname and device path are used to construct a device name that follows the existing remote device syntax (`rd=hostname:device_path`).

Host Device Selection for Recoveries

The Solstice Backup server contains an attribute in the Jukebox resource for designating a host for use in recoveries when the host controlling a required device is not already listed in the Client resource's Storage Nodes attribute. This attribute is the Read Hostname attribute, and its default value is the Solstice Backup host that is controlling the first device in the shared autochanger. However, you can change the host to another value contained in the autochanger.

This new functionality designates an alternative device to use for recoveries and cloning if a storage node's preferred hosts are unavailable. Normally, when a recovery is requested, the Solstice Backup server will match the host that is controlling the device required for the recovery with the hosts listed in the Storage

Nodes list located in the Client resource. After the first match, the recovery or clone operation takes place. If there is no match, then the request reverts to the host listed in the Read Hostname attribute.

Example: Data Recovery Using a Shared Autochanger

FIGURE 10-3 displays an example of a Solstice Backup datazone environment that consists of a server venus, a client earth, and a storage node mars. There is a shared autochanger with the first device connected to venus, and a second device connected to mars.
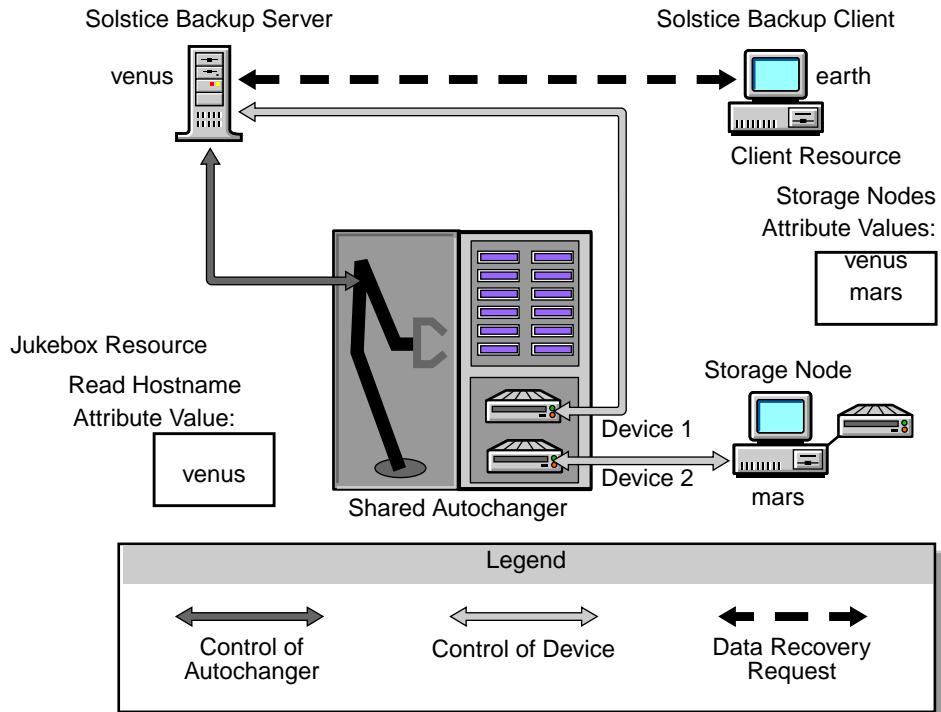
FIGURE 10-3 Data Recovery Using a Shared Autochanger

The client earth requests a data recovery from a volume in the shared autochanger. The needed volume is not currently mounted in any device. To initiate the recovery, the server venus searches the hostname entries listed in the Storage Nodes attribute of earths's Client resource for a match with one of the hostname devices in the autochanger. A hostname match is made for storage node mars. The needed volume is mounted on that device, and the recovery occurs.

If a hostname match was not found, then venus would refer to the hostname value listed in the Read Hostname attribute of the Jukebox resource. By default, this value is venus. The needed volume is then mounted in the device controlled by venus, and the recovery occurs.

Host Device Selection for Cloning

When a cloning request is received, the Solstice Backup server uses the host listed in the Read Hostname attribute of the Jukebox resource to read the volume being cloned. If the required volume is already mounted in a device, then the Solstice Backup server accesses that particular device.

Labeling, Loading, and Unloading Volumes on Shared Autochanger Devices

Autochanger sharing also affects tasks related to volumes present in autochanger devices, such as:

■ Labeling volumes

■ Loading and unloading volumes for labeling, client backups, and data recoveries

The Solstice Backup command to manage autochanger volumes, nsrjb, has been modified to allow management of volumes on shared autochanger devices controlled by other hosts. The nsrjb command must always be run from the Solstice Backup host controlling the autochanger (such as having control of the robotic arm) regardless of whether the autochanger is shared or not.

Note – When you use the nsrjb command with the -f *devicename* option and your device name contains backslashes, you must use the backslash "escape" character with the backslash. Therefore, a device with the name rd=sn1:\\.\Tape0 would be entered at the command line as rd=sn1:\\\\.\\Tape0.

Managing Volumes on Server Devices

For devices controlled by the Solstice Backup server, you manage autochanger volumes using the nsrjb -f command with the device path, even if the jbconfig command for the shared autochanger was run from another host. For example, if you wanted to load a volume into a shared autochanger device that is controlled by the Solstice Backup server, you would enter the following command at the prompt:

```
# nsrjb -l -f device_path volume_name
```

Managing Volumes on Storage Node Devices

When a shared autochanger is defined within the datazone, you must specify the full remote device syntax name when using the nsrjb -f command for devices controlled by a storage node:

```
# nsrjb -l -f rd=host:device_path volume_name
```

where host is the hostname of the appropriate node that controls the device.

Note – In a datazone where no shared autochangers reside, the Solstice Backup server automatically adds the remote device host (rd=host) information; therefore, you only need to input the local device syntax (device_path). However, if there is one or more shared autochangers in the datazone, then you must always enter remote device syntax (rd=host:device_path) information with the device path.

Backups and Recoveries on a Shared Autochanger Device

Autochanger sharing also requires that the host of a volume being selected for a backup or recovery be further specified to reflect a particular host. Normally, the Solstice Backup server automatically performs this activity and requires no user input. However, in cases where you need to perform this task at the command line, the -J option has been added to nsrjb. The -J option is only available when specified with one of the following:

- The nsrjb -l option for loading a volume
- The nsrjb -L option for labeling a volume

It is not available for use in conjunction with the `nsrjb -f` option for specifying a device path.

Automatic Unmounting of Volumes

At times, a volume that is mounted in one device may be needed by another device on the same autochanger. For example, data being recovered by one device may span more than one volume, and the required volume is mounted on another device. To address this need, you can define a value in the Idle Device Timeout attribute in the Solstice Backup resource for that autochanger. The Idle Device Timeout attribute specifies the number of minutes a mounted volume can remain idle before it is automatically unmounted from the device and returned to its slot, where it can then be accessed by another device.

Setting Compatible Block Size for UNIX and Microsoft Windows

There are different SCSI hardware limitations between UNIX and Microsoft Windows platforms, which can lead to compatibility problems. For example, if you have a device defined in UNIX and it is physically attached to a Windows host bus adapter (HBA), it is possible to define a block size greater than that allowed by the Windows hardware. This could lead to I/O errors in both write and read states on this device. If you are going to use both platforms, you must determine a common block size acceptable to both.

Determining the Problem

The `scanner` program can be used to determine the allowed block size. Enter the following command at the prompt:

```
# scanner -v device_name
```

where *device_name* is the name of the device as defined for the server. The block size is included in the following output from this command:

```
scanner: volume id 2641870849 record size 98304
```

The record size entry is the block size of the device, in bytes.

Solutions

To solve problems with block size compatibility:

■ On UNIX, you can change the block size by setting the following environment variable to the greatest common value for both systems. For example:

```
setenv NSR_DEV_BLOCK_SIZE_DEVICE_NAME "value"
```

where *DEVICE_NAME* is the backup device type available to your Solstice Backup server and *value* must be a multiple of 32 KB, with a minimum value of 32 KB. Note that the syntax for *DEVICE_NAME* is all upper case, with underscores (_) replacing blank spaces. Therefore, a device displayed in the Solstice Backup software as "4mm 20gb" would be listed as follows:

```
4MM_20GB
```

■ On Microsoft Windows, you can install a later model HBA or upgrade to drivers that can support up to 128 KB blocks for the previous UNIX example. Windows also accepts the same environmental variable format for setting block size as listed for UNIX.

Autochanger Operations

This section includes the following information for performing autochanger operations:

You may need to perform some of these tasks on a daily basis and others less frequently. The frequency with which you perform these tasks depends on:

■ The number of volumes you label at one time
■ How often you change the contents of your autochanger
■ How often you need to mount volumes for manual backups

▼ To Designate Available Autochanger Slots

The available slots feature allows you to control which volumes the Solstice Backup server uses for backing up data. The server uses all of the volumes in the autochanger to perform recoveries. However, you can control the volumes automatically selected for backups by designating a range of available slots in the autochanger.

For example, you may designate slots 1 through 5 for your Solstice Backup backups in an autochanger that contains 10 slots. The entries can be a range of slot numbers or a single slot number. The range of available slots you enter is the same for all of the storage devices in the autochanger. You cannot enter a different range for each device located in an autochanger.

> **Caution –** Make sure you have placed volumes in all the available slots of the autochanger so the Solstice Backup server can proceed uninterrupted with an automatic backup.

To enter the number of available slots:

1. In the Solstice Backup Administrator program, select Jukeboxes from the Media menu to open the Jukeboxes window.

2. Select the autochanger whose available slots you want to change in the Jukeboxes attribute.

3. Enter the range of available slots into the Available Slots attribute.

4. Click Apply.

Note that with two-sided media, the number of available slots is effectively doubled. For example, if you have 32 optical disks, labeled "atlas.001.a" to "atlas.032.b," you have a total of 64 sides, and therefore 64 slots from which to choose.

▼ To Clean Autochanger Devices

Maintaining a storage device in good working order requires periodic cleaning. The Solstice Backup server provides automatic cleaning of devices located in an autochanger. The server does not support automatic cleaning for stand-alone devices. Cleaning autochanger devices requires you to configure attributes in both the Jukebox and Device resources of each appropriate device.

To clean autochanger devices:

1. In the Solstice Backup Administrator program, select Jukeboxes from the Media menu to open the Jukeboxes window.

2. Select the autochanger for which you want to configure cleaning.

3. Enter the slot number of the autochanger that contains the cleaning cartridge in the Cleaning Slots attribute. If more than one slot contains a cleaning cartridge, specify the range with a hyphen, for example, "8-9."

4. Enter the number of times the cleaning cartridge can be used before it needs to be replaced in the Default Cleanings attribute.

5. Select Yes in the Auto Clean attribute to enable autocleaning.

6. Click Apply.

7. Select Media>Devices to open the Devices window.

8. Select the appropriate device in the Devices attribute.

9. Select Yes in the Cleaning Required attribute.

10. Enter a value in the Cleaning Interval attribute. This attribute indicates the amount of time between cleanings. You can specify the time in days, weeks, or months.

11. Click Apply.

   Note that the last time the device was cleaned is listed in the Date Last Cleaned attribute.

12. Repeat these steps for each device included in the autochanger.

To ensure there is no interference with other autochanger operations, the Solstice Backup server only cleans devices before mounting or after unmounting a volume in a device. For recommendations on the frequency and cleaning methods for your autochanger devices, check the autochanger manufacturer's documentation.

TABLE 10-3 lists the notification messages the Solstice Backup software displays to inform you about cleaning cartridge operations.

TABLE 10-3   Notification Messages

Notification Message	Meaning
Device cleaning required	Auto Clean is disabled and the device needs to be cleaned.
Device cleaned	Auto Clean is enabled and the device has been cleaned.
Cleaning cartridge required	Auto Clean is enabled and there are no usable cleaning cartridges available.
Cleaning cartridge expired	The cleaning cartridge has been used the specified number of times.

▼ To Use a Nondefault Slot for the Cleaning Cartridge

To use a nondefault slot for the cleaning cartridge:

1. Insert the cleaning cartridge in the slot you want.

2. Select Media>Jukeboxes to open the Jukeboxes window.

3. Select View>Details to display hidden attributes.

4. Select Yes in the Auto Clean attribute.

5. Specify the slot number you want to use for the cleaning cartridge in the Cleaning Slots attribute.

6. The Default Cleanings attribute lists the number of uses left on the cleaning cartridge. Record this number for use in Step 8.

7. Change the Available Slots attribute to reflect the range of slots available for data volumes.

   If the slot for your cleaning cartridge is not the first or last slot in the autochanger, you must specify two ranges of slots in the Available Slots attribute. This is because the inventory of the autochanger must be performed in two steps, once for each range of slots.

   For example, if the autochanger contains 11 slots and slot 6 is used for the cleaning slot, specify:

   ```
   1-5

   7-11
   ```

8. Enter the following command at the prompt:

   ```
   # nsrjb -U uses -j autochanger -S slot
   ```

   - Replace *uses* with the number of uses left on the Default Cleanings attribute (the number you wrote down in Step 6).
   - Replace *slot* with the slot you now use as the cleaning slot.
   - You can omit the -j option if there is only one autochanger.

   Every time you replace the cleaning cartridge in the autochanger, you must run the nsrjb program to specify the number of uses left.

⚠️ **Caution –** If your autochanger does not support the bar code labeling option, you must issue the command shown in Step 8 to tell the autochanger to add the cleaning cartridge to its inventory.

▼ To Replace Cleaning Cartridges

When a cleaning cartridge expires, the Solstice Backup software displays the following message:

```
Date server_name: Solstice Backup Cleaning cartridge expired:
(notice) Cleaning cartridge required in Jukebox 'autochanger_name'
```

Replace the expired cartridge and enter the following command at the prompt:

```
# nsrjb -U number_of_cleanings
```

If you do not enter the nsrjb -U command along with the number of cleanings, the Solstice Backup software does not recognize the new cartridge and continues to send the "cleaning cartridge required" notification.

If your autochanger does not support bar code labeling, you must enter the following command at the prompt to tell the autochanger to add a cleaning cartridge to its inventory:

```
# nsrjb -U number_of_uses -j autochanger_name -S slot
```

– Replace *number_of_uses* with the number of cleanings the cartridge supports.
– Replace *autochanger_name* with the name shown in the Jukebox attribute.
– Replace *slot* with the slot number where you inserted the cleaning cartridge.

Auto Media Management with Autochangers

Auto media management further automates the backup procedure by freeing administrators from the task of labeling volumes before backup. When auto media management is enabled in the Jukebox resource, the Solstice Backup software assumes that all the volumes loaded in the autochanger may be handled exclusively as Solstice Backup media. Volumes that appear unlabeled are considered blank and

are automatically labeled, mounted, and overwritten with new data. The Solstice Backup software only uses volumes that it considers unlabeled if it cannot locate a writable volume or a volume eligible for automatic recycling.

The Solstice Backup software does not recognize the following volume labels and considers these volumes unlabeled:

- Volumes without a label

- Volumes labeled with something other than a Solstice Backup label

- Volumes with Solstice Backup labels that were written in a density different from the device in which the volume is currently loaded

  For example, if you performed Solstice Backup backups on a new 8 mm tape drive, an older 8 mm tape drive might not be able to read the volumes from the newer device because it used a different density to write the data.

Caution – Be careful about sharing Solstice Backup volumes between different storage devices if you have auto media management enabled. You could potentially relabel and overwrite valuable data.

The auto media management feature is located in both the Jukebox and Device resources. For autochangers, you only need to enable auto media management in the Jukebox resource. The Solstice Backup server does not allow you to enable auto media management in the Device resource if the device is located in an autochanger.

If you back up to a stand-alone device and want to use the auto media management feature, see "Auto Media Management" on page 235.

When auto media management is enabled, the Solstice Backup server assumes that the autochanger is exclusive to the server, and the volumes loaded in the device are handled as Solstice Backup media. Do not enable this feature if the autochanger or device is a shared resource. The Solstice Backup server assumes it can label and overwrite any media it considers unlabeled.

However, if the autochanger is a shared device, you may use the available slot range to restrict the volumes used for backup. If you place or store other volumes in the autochanger, you do not want to use slots that are not included in the available slot range.

How the Solstice Backup Server Selects a Volume for Writing

If auto media management is enabled and the Solstice Backup software is connected to an autochanger, the autochanger searches for a mounted volume for writing. If no mounted volume is a candidate for writing, the Solstice Backup software looks for an unmounted, appendable volume from the appropriate pool that can be mounted.

If none is found, the Solstice Backup software looks for a recyclable volume from the required pool that can be mounted. If a recyclable volume is found, the software mounts the volume and relabels the volume. As part of the relabeling operation, the Solstice Backup software removes all entries for the save sets and volume that remain in the media database. A new entry for the relabeled volume is added to the media database.

The Solstice Backup software formats the label by obtaining the next volume name from the pool's associated label template. It automatically creates a label that incorporates the next-available-unused label, which is not necessarily the next-in-sequence label, depending on previous labeling activity or volume deletions. Administrators are advised to affix a plain language label on the outside of the volume as well for easier identification.

When a volume is recycled, the expiration date is automatically reset to the value of the Volume Expiration attribute of the Device resource (two years by default). If an administrator uses the nsrjb program to relabel a volume in an autochanger, an expiration date can be specified from the command line.

If no recyclable volumes from the required pool are found, the Solstice Backup software looks for an unlabeled volume. If any unlabeled volume is found, it mounts the new volume and labels it. An entry for the newly labeled volume is added to the media database.

Mounting Autochanger Volumes

Typically, autochangers automatically mount volumes needed for backup and recovery if they are located in the autochanger. When a volume located in an autochanger is required for backup, Solstice Backup software uses the following hierarchy to select a volume from the appropriate pools that are not currently in use:

- Already mounted, appendable volume
- Unmounted, appendable volume
- Unmounted, recyclable volume
- Unmounted, unlabeled volume (autochanger must have auto media management enabled)

▼ To Mount Autochanger Volumes

To mount autochanger volumes:

1. At the main Solstice Backup Administrator program window, select the appropriate autochanger in the Devices window.

2. Select Media>Mount to open the Jukebox Mounting window.

3. Select a volume from the Slot attribute and click OK.

The Solstice Backup server automatically mounts the volume without any further prompting. The Devices attribute in the Solstice Backup Administrator program reflects the device's mounted status.

It is a good idea to place an adhesive label on the autochanger identifying its device pathname(s). This reminds you which device pathnames belong to the autochanger. This is especially helpful when you back up to more than one autochanger.

▼ To Mount Uninventoried Volumes

This section describes how to mount a volume that is not included in the autochanger inventory, but which is a valid (properly labeled) Solstice Backup volume.

To mount an uninventoried volume:

1. At the main Solstice Backup Administrator program window, select the appropriate autochanger in the Devices window.

2. Select Media>Inventory to open the Jukebox Inventory window.

3. Inventory the slot by entering the same slot number in the First Slot and Last Slot attributes. For more information on inventorying, see "Inventorying Autochanger Volumes" on page 400.

   The Solstice Backup software will not inventory the slot you specified. Now you can mount the inventoried volume as described in "To Mount Autochanger Volumes" on page 393.

Note – Unlabeled media is not considered "valid" for mounting in this case because media of this type can only be mounted for labeling, not inventorying. An I/O error will occur and the volume will be ejected if you try to mount uninventoried volumes using unlabeled media.

▼ To Unmount Volumes

To unmount volumes:

1. At the main Solstice Backup Administrator program window, select the appropriate autochanger in the Devices window.

2. Select Media>Unmount to unmount the volume.

3. Select the device with the volume you want to unmount from the Devices attribute and click Unmount.

The Solstice Backup server automatically unmounts the volume without any further prompting and returns the volume to its appropriate slot. The Devices attribute in the Solstice Backup Administrator program reflects the appropriate device's unmounted status.

▼ To Label Autochanger Volumes

The Solstice Backup server uses the label template selected in the Pool resource to label the volumes loaded in the autochanger. Before starting a labeling session, select a label template from the Pool resource, or create your own. For more information on label templates, see "How the Solstice Backup Server Uses Label Templates" on page 208.

It will take the Solstice Backup server some time to label the volumes in the autochanger. You may want to perform the labeling process at a time when you do not need to back up or recover files.

To label an autochanger volume:

1. At the main Solstice Backup Administrator program window, select the appropriate autochanger in the Devices window.

2. Select Media>Label to open the Jukebox Labeling window.

3. Enter a label in the Starting With attribute.

   You can enter your own label only if you are labeling a single volume in the autochanger. If you are labeling a range of volumes, the label is automatically assigned based on the pool you select.

4. Complete the First Slot and Last Slot attributes:

   - To label a range of slots, enter the range in the First Slot and Last Slot attributes.
   - To label only one volume, enter the same number in both the First Slot and Last Slot attributes.

5. Select a pool in the Pool attribute for the volumes being labeled.

   The Solstice Backup server uses the label template associated with the selected pool to label the volumes. Remember, only pools that have been enabled in the Pool resource will appear in the Jukebox Labeling window.

6. Select the Manual Recycle attribute only if you want to recycle the volume manually.

> **Note –** Use this option if you want to keep the data beyond the retention policies associated with the data and protect the volume from automatic recycling. When you no longer want to keep the data, you must change the volume mode to either *recyclable* or *appendable*.

7. Click OK.

When a valid Solstice Backup label already exists on the media you are trying to label, a confirmation box appears asking you if you really want to relabel the volume. The confirmation is important because it keeps you from accidentally relabeling the media. When a volume is relabeled, its contents under the previous label cannot be recovered.

> **Caution –** If you label a specific range of volumes, the name that appears in the Volume Label must match the label template. If you label a single volume, you can use any name; it does not have to match the label template.

Labeling Volumes with Bar Codes

Labeling media with bar code labels provides two distinct advantages over non-bar code labeled volumes. Bar code labeling:

- Significantly accelerates volume inventory.
- Improves labeling accuracy.

Inventorying with bar code labels is fast and efficient because you do not have to mount the volumes into the device. Instead, the autochanger scans the external bar code labels with an infrared light while the volumes remain in their slots. Inventorying with bar code labels greatly reduces the time it takes to locate a volume or determine the contents of an autochanger.

Bar code labels also provide greater labeling accuracy because the labels are adhered to the media prior to being loaded and scanned in the autochanger. Once the autochanger scans the bar code, the Solstice Backup server records and tracks the label in the media database.

The Solstice Backup server uses bar code labels only to inventory volumes. It is important to update and maintain an accurate account of the autochanger contents by using the `inventory` command. Otherwise, locating a volume for backup or recovery (whether or not it uses bar code labels) is impossible because the Solstice Backup server does not know which slot contains the needed volume.

A volume must have a label, but it does not require a bar code label.

Caution – The autochanger includes hardware to read the bar code labels. The bar code information is then forwarded to the Solstice Backup server. Problems reading bar code labels indicate hardware problems. In the event of a bar code-related problem, consult your autochanger documentation or hardware vendor.

▼ To Label a Volume with Bar Codes

To label a Volume with bar codes:

1. Adhere the bar code labels to your volumes.

2. Place the volumes with the bar code labels in the autochanger.

3. Select the appropriate autochanger device in the Devices attribute of the Solstice Backup Administrator program.

4. Select Media>Jukeboxes to open the Jukeboxes window.

5. In the Jukeboxes attribute, select the autochanger whose volumes you want to label with bar codes.

6. Select Yes in the Bar Code Reader attribute.

   If you want the volume label to match the bar code label, select Yes in the Match Bar Code Labels attribute.

   In either case, you will have two labels: the volume label (which may or may not match the bar code label, depending upon your bar code choices) and the bar code label.

7. Click Apply.

8. Select Media>Label to open the Jukebox Labeling window.

   If you enabled the Match Bar Code Labels attribute in the Jukebox resource, the Volume Label attribute displays the following label:

```
Bar code label
```

   Otherwise, the Volume Label attribute displays a label template name.

After you have selected your bar code label choices and labeled the volume, both the volume label name and bar code label for the volume appears in the Volumes window.

Bar Code Labeling Tips

The volume label and bar code label are both pieces of information that the Solstice Backup server uses to identify volumes. Both labels are recorded in the media index. However, the volume label is also recorded internally on the media (internal volume label). The Solstice Backup server uses bar code labels to inventory volumes and uses the volume labels to identify the volumes needed for backup and recovery. You can require that the volume label match the bar code label by making the appropriate selections for bar code labeling in the Jukebox resource.

When Solstice Backup software relabels volumes automatically, it reuses the original volume label name. You can only change a label name if you relabel the volume manually. The Solstice Backup software scans the bar code label during the labeling process and updates the media index with the new volume name and its associated bar code label.

- Do not use identical bar code labels for any of your Solstice Backup volumes. Using identical labels defeats the purpose of using bar code labels, which is to speed up the inventory process and ensure labeling accuracy.

- Give every volume a unique volume label. If you try to label a second volume with an identical bar code label and you have enabled Match Bar Code Labels attribute in the Jukebox resource, the Solstice Backup server displays an error message and will not allow you to label the second volume. The error message identifies the autochanger slots containing the two volumes with identical labels and the bar code label.

  To correct the problem, you must either apply a different label to one of the volumes and begin the labeling process again, or you must disable Match bar code labels in the Jukebox resource while labeling the second volume.

- It is not necessary to label existing volumes with bar code labels if they are stored in a vault or offsite for long periods at a time. This is because you will not inventory these volumes often, if ever.

- Before using bar code labels on your existing volumes, first adhere the bar code labels to them. Then, load and mount each volume individually, so the Solstice Backup server can match the bar code label with the existing volume label.

- You may purchase a variety of bar code labels from a third-party vendor. You can choose numeric labels, alphanumeric labels, or a special combination of numbers and characters to meet your labeling needs. Furthermore, you can order bar code labels that match your current volume labeling scheme.

- Use a consistent labeling scheme. If you label your volumes with the server name and an extension such as "001," order a range of labels starting with "*server_name*.001" and ending with "*server_name*.100" or as wide a range as you need. Labeling instructions for bar code labels should be provided with the autochanger hardware documentation. If you have questions about bar code labels, call the hardware manufacturer.

Using a consistent labeling scheme helps you better organize and track your volumes. It also aids the inventory process if all the volumes use bar code labels rather than just a limited number of them.

▼ To Use Unmatching Volume and Bar Code Labels

Note – If you choose not to match the volume label to the bar code label, you should attach labels to the outside of your media.

To use unmatching volume and bar code labels:

1. Apply the bar code labels to your volumes.

2. Place the volumes with the bar code labels in the autochanger.

3. Select the appropriate autochanger device in the Devices attribute of the Solstice Backup Administrator program.

4. Select Media>Jukeboxes to open the Jukeboxes window.

5. In the Jukeboxes attribute, select the autochanger whose volumes you want to label with bar codes.

6. Select Yes in the Barcode Reader attribute.

7. Select No in the Match Bar Code Labels attribute.

8. Click OK.

    The Solstice Backup server uses the next available label from the label template for the volume name. It labels the volumes and records both labels in the media index.

9. Select Media>Inventory to open the Jukebox Inventory window.

10. Click OK.

11. Select Media>Volumes or click the Volumes toolbar button to open the Volumes window.

12. Determine the volume label and bar code label for each volume. Use the information in the Volumes window to match the correct volume labels to the bar code labels. This is the easiest way to make sure that you adhere the correct volume labels to the volumes with the existing bar code labels.

If you enabled the bar code functionality and did not adhere a bar code label to the volume, you will receive an error message informing you that a bar code label does not exist.

Inventorying Autochanger Volumes

When Solstice Backup software labels the contents of an autochanger, it registers the location of the volumes in the autochanger slots when it assigns the volume label. *This is called taking inventory. When you* inventory the volumes in the autochanger, Solstice Backup software reads the label of each volume and records its slot number. If you do not change the volumes in the autochanger after labeling them, the Solstice Backup server can access the volumes because each volume label is assigned to a specific slot.

However, if you change the contents of the autochanger without performing the labeling process, or if you move volumes into new slots, you must inform the Solstice Backup software that the autochanger now holds a different set of labeled volumes or that the volumes are in a different order. For example, if you have more than one magazine for an autochanger, you must inventory each time you remove one magazine and load another one into the autochanger.

When you label the volumes in a new magazine, you do not need to inventory them because the Solstice Backup software automatically records the slot number in which the newly labeled volume is located.

The Solstice Backup software provides bar code labels for speeding up the inventory process. Use bar code labels if you have a large number of volumes or change your autochanger contents often. For more information, see "Sharing Autochangers between Solstice Backup Hosts" on page 379.

▼ To Inventory Autochanger Volumes

To inventory autochanger volumes:

1. Load the new volumes in the autochanger.

2. At the main Solstice Backup Administrator program window, select the appropriate autochanger in the Devices window.

3. Select Media>Jukeboxes to open the Jukeboxes window.

4. Enter the range of the slots you want to inventory in the First Slot and Last Slot attributes, or leave the default values as is to inventory all the slots.

5. Click OK.

If your volumes do not have bar code labels, the Solstice Backup software must mount each volume, read its label and unmount it. In this case, the inventory process may take some time to complete. After completing an inventory, the Solstice Backup Administrator program displays the results in the Volumes window.

I/O Errors

If you inventory volumes in an autochanger, and a volume is not labeled, the following message appears:

```
volume tape label read, I/O error
```

The process by which an inventory attempts to verify the label for each volume in an autochanger cannot distinguish between an I/O error caused by reading an unlabeled volume and an I/O error caused by reading a bad volume. Once the volume is loaded into the device, the nsrjb program makes a request to nsrd to load the volume and verify its label. The nsrd program allocates an nsrmmd session to perform the verify label operation; nsrmmd then attempts to position the volume to read the label and generates the error.

Do not run an inventory command on an autochanger filled with new unlabeled volumes. Instead, enter the following hardware reset command at the prompt:

```
# nsrjb -HE
```

The nsrjb -HE command (hardware reset with initialize element status) synchronizes the state of the library, after which you can label the volumes in the autochanger. At this point, you can run the inventory on the volumes in the autochanger, using either the steps described in this section, or by running nsrjb - I.

Using Autochangers with a Volume Import/Export Capability

The Solstice Backup server supports the use of the SCSI-II import/export feature found in many brands of autochangers. Depending on the autochanger model, this feature is also known as cartridge access port (CAP), mail slot, and loading port. The import/export feature deposits and withdraws volumes from slots in the autochanger. This feature enables the operator to deposit and withdraw cartridges

without invalidating the device inventory list. Normally, if the operator opens the door to load or unload media, the element status of the autoloader is invalidated, requiring the long operation of reinitializing the autochanger. However, the Solstice Backup server does not automatically inventory the volume after a deposit and withdraw.

Note – The procedures in this section do not apply to silos. For more information, see "Using Silos with a Volume Import/Export Capability" on page 414.

▼ To Deposit Volumes

To deposit volumes:

1. Log in as root on the Solstice Backup server or storage node computer.

2. Press the button on the front of the autochanger to move the cartridge holder forward and open the CAP.

3. Place the volume in the holder.

4. Press the button again to move the media into the autochanger and close the CAP.

5. Enter the following command at the prompt:

```
# nsrjb -d -S slot
```

where *slot* is the slot number for the volume.

If you do not specify a volume name with the nsrjb command, the Solstice Backup server marks the slot with a "-*" to indicate that the volume in that slot is unknown. Inventory the slot with the following command:

```
# nsrjb -I -S slot
```

6. To verify that the volume was deposited in the correct slot, mount the volume.

▼ To Withdraw Volumes

To withdraw volumes:

1. Log in as root on the Solstice Backup server or storage node computer.

2. Enter the following command at the prompt:

```
# nsrjb -w -S slot -P port
```

   where *slot* is the slot number for the volume and *port* is the value assigned to the control port.

3. When the volume reaches the CAP, press the button to open the CAP.

4. Remove the volume and close the CAP.

Tips and Suggestions for Using Autochangers

This section provides additional suggestions to help you use your autochanger and backup media effectively and reliably. The following information is provided:

Autochanger Notifications

The Solstice Backup server uses notifications to send messages about Solstice Backup events. A number of preconfigured notifications are included to inform you of the following situations:

- Volumes in the autochanger are 90 percent full
- Autochanger needs more volumes to continue

- Autochanger has a mechanical problem
- Autochanger device needs cleaning
- Cleaning cartridge needs attention

The Notifications resource displays all of the notifications, including:

- Three tape mount requests
- Two cleaning cartridge notifications
- Two device cleaning notifications

For more information, see "Notifications" on page 99.

The Solstice Backup software automatically mounts a required volume as long as the volume is loaded in the autochanger. If the recovery requires a volume that is not loaded in the autochanger, the *Tape mount request 1* notification sends a message to the Solstice Backup Administrator program, which informs you to mount a specific volume.

After you correct an autochanger problem, you might need to mount a volume so the Solstice Backup server can continue to back up or recover files. Check the Solstice Backup Administrator program for messages.

Recycling Compared to Adding More Volumes

The Solstice Backup server saves files on volumes marked *appen* (appendable). If the volumes inside the autochanger are marked full, they cannot receive backups.

You can do one of the following tasks with volumes marked *full*:

- Remove the full volumes and replace them with new media if you need to keep the volumes for long-term storage.
- Change the volume mode to *recyc* (recyclable) if you do not need the data on the full volumes. The Solstice Backup server will overwrite the data with new backups, yet maintain the existing labels. This is the only instance in which you do not need to relabel a volume to make it eligible to be overwritten with new data.

When all of the save sets on the volume have passed the time period specified by the retention policy, the mode of a volume is automatically changed to recyclable.

There are advantages to both recycling media and adding more media to a pool. By recycling, you reuse the same volumes and do not add new volumes to the pool. However, the media may wear out over time and have a higher failure rate.

On the other hand, if your site requires that you maintain the backups in storage for a certain period of time, then you may have to add more media to the pool instead of recycling it. For example, an autochanger may need new volumes every three

months if the company policy is to maintain the backups for a year. In this case, you need to keep adding new media to the pool until you can recycle the volumes that contain expired or old backups.

Using Pools with Autochangers

If you plan to have full and nonfull backups, we recommend that you estimate the number of volumes needed for your full backups and assign them to the Full pool. This ensures that your full backups are located in a consecutive range of slots in the autochanger. This allows you to remove all of the volumes at the same time.

Device Calibration

For information about the frequency and method for calibrating the loading mechanism for the device, refer to the autochanger manufacturer's documentation.

▼ To Disable the Element Status

Some autochangers support the element status feature. The Solstice Backup software uses the element status feature to determine whether a slot contains a volume. If volumes have bar codes, the element status feature speeds up the inventory process by eliminating the need to load and read volume labels. If the volume has a bar code label, Solstice Backup software reads it without loading the volume.

To disable the element status option for your autochanger if your volumes do not have bar codes:

1. In the Solstice Backup Administrator program, select Jukeboxes from the Media menu to open the Jukeboxes window.

2. Select View>Details to display hidden attributes.

3. Deselect the Elements_Status option in the Jukebox Features attribute.

4. Click Apply. The element status option is now disabled.

▼ To Reset the Autochanger

You might need to reset an autochanger if it is unable to complete an operation, or if it appears to be in a "confused" state. For example, the Solstice Backup server acts as if a slot is empty whereas in reality a volume is loaded in the device. The opposite might be also be true: the autochanger acts as if the slot is full, but in reality the slot is empty. Using the reset feature can return an autochanger with these characteristics back to normal functionality.

To reset the autochanger:

1. Enter the following command at the prompt:

```
# nsrjb
```

Make a note of the output. This command displays information about the volumes loaded in the devices, the volume labels, and the slots from which the volumes were loaded into the devices.

2. Manually unload the volume from the device. For information on how to do this, refer to your autochanger hardware documentation.

3. Manually reinsert the volume into the cartridge access port (CAP), flipping it over so that the "A" side faces up.

4. Reload the volume from the CAP into the slot from which it originally came, using the control panel on the autochanger.

5. Enter the following command at the prompt:

```
# nsrjb -H
```

The autochanger will now reset itself.

Optical Autochangers

If your optical autochanger does not work with the Solstice Backup software, your operating system might not currently include support for the optical autochanger media device.

When you create the Solstice Backup device names for an optical autochanger, you should use the raw name of the device.

If you experience a power outage, the optical autochanger loses track of the state of its backup volumes. Issuing the nsrjb -H command might not successfully reset the autochanger. If you see output similar to the following, it means the nsrjb -H command failed to reset the autochanger:

```
nsrjb: Autochanger error, Invert operation not supported
```

Silos

The Solstice Backup software manages silos and silo media similarly to autochangers. For best results, you should familiarize yourself with autochanger operations before moving on to silos.

A silo tape library (STL) is a peripheral that usually contains many storage devices. Silos are controlled by silo management software, which is provided by the silo vendor and installed on a silo server. The silo server cannot be the same computer as the Solstice Backup server.

The silo can be shared among many applications, systems, and platforms. Like autochangers, silos make data and media operations more automatic. Silos can load, change, and manage volumes, and clean the devices automatically.

How Solstice Backup Software Interacts with a Silo

A Solstice Backup server acts as a client of the silo management software, which resides on the silo server. The Solstice Backup server communicates with the silo through the Silo Tape Library Interface (STLI), which must be installed on the Solstice Backup server that uses the silo.

To access the volumes and devices in a silo, the Solstice Backup server sends a request to the silo management software, in the form of an STLI call. For example, to mount a volume in a silo device, the Solstice Backup media service sends a request to the silo management software to mount the volume into a particular device in the silo. The silo server responds to the request and mounts the volume in the requested device.

The silo management software controls many of the operations that Solstice Backup software controls with an autochanger. For example, the silo management software keeps track of the slot where each silo volume resides, and may control the deposit and withdrawal of volumes and automated cleaning of silo devices.

▼ To Install Silos

To install a silo for use with Solstice Backup software:

1. Install the silo management software on the silo server.

2. Install the STLI on the Solstice Backup server (if necessary).

   Do not install the STLI library on the following models, because all the necessary software was installed when you installed the Solstice Backup software:

   - IBM 3494 on Solaris and AIX
   - StorageTek on Solaris, AIX, HP-UX, and DYNIX/ptx
   - Advanced Digital Information Corporation (ADIC) EMASS/Grau on Solaris, AIX, HP-UX, and Windows NT

   For other silo models, install the STLI library on the Solstice Backup server or storage node that uses the silo. Follow the instructions provided by the silo vendor.

3. Ensure that the Solstice Backup server is properly connected to the media devices in the silo.

4. Run the jbconfig program to configure the devices in the silo.

   The jbconfig program prompts you to enter the following information:

   - Device type to install. Enter **3** (for STLI Silo).
   - Type of silo.
   - Name of the silo. Enter any alphanumeric string.
   - Hostname of silo server.
   - Whether to enable automatic device cleaning. Enter **no**. The silo server controls device cleaning in silos.
   - Pathname of the STLI software. Enter the full path on the Solstice Backup server where the STLI software resides.
   - Number of devices to configure.
   - For each device you configure, the program prompts you for the pathname of the device, media type of the device (if the device is not already configured), and silo name of the device.

   For additional information, see "To Add Autochangers" on page 361.

5. License the Silo Support Module. For general information on licensing, the Solstice Backup Roadmap.

Silo Device Naming Conventions

The `jbconfig` program prompts you for the silo name of the storage devices. The silo name is the name that the silo management software uses to refer to the storage device. Depending on the type of silo you have, the device name can take several forms. This section describes the naming conventions of the currently supported silos.

StorageTek

The StorageTek (STK) silo management software uses either a program called ACSLS that runs on a UNIX system, or a program called Library Attach that runs on an MVS system. These program name devices according to a coordinate system based on the physical location of the devices in the silo.

For tape drives, the name consists of four digits separated by commas:

- The first digit refers to the automated cartridge system (ACS) with which the drive is associated.
- The second digit refers to the library storage module (LSM) in which the drive is located.
- The third and fourth digits refer to the panel and slot location in which the drive is located.

A typical name for an STK drive looks like 1,0,1,0.

You must ask the silo administrator for the drive names of the devices that the Solstice Backup server can use. There is no method to find this information from the Solstice Backup server. To connect to more than one drive, find out the SCSI IDs for each drive and properly match the IDs to the silo names. If you accidentally swap operating system device names and silo names, you can only mount and unmount volumes; you cannot read or write to the volumes after they are mounted. To reconfigure the device names properly, use the Solstice Backup Administrator program to change the order of the device names in the STL Device Names attribute of the Jukeboxes resource.

IBM 3494

The silo management software for the IBM 3494 names devices with an eight-digit number to identify the 3590 drives in the silo. Use the appropriate utility to obtain the device names, as follows:

- On an AIX system, the `jbconfig` program gets the name of the device from the device driver and displays the device name as the default value.

- On a Solaris system, you must use the IBM supplied `mtlib` command (`mtlib -l` *library_name* `-D`) to determine the names of all the devices in the 3494. Either ask the silo administrator which device is reserved for the Solstice Backup software, or test to decide which silo drive name matches with each Solaris device name.

EMASS/Grau

The silo management software for the EMASS (in North America) or Grau silos is a program called DAS. DAS acts as a front end for the silo control program called AMU. When the silo is configured, the silo administrator gives each drive a symbolic name. The symbolic name can be any alphanumeric string.

▼ To Set up DAS to Work with Solstice Backup Software

To set up DAS to work with Solstice Backup software:

1. Ask the silo administrator to configure DAS to accept commands from your Solstice Backup server or storage node computer.

2. Ask the silo administrator to do one of the following:
   - Use the `dasadmin allocd` command to allocate one or more devices to the Solstice Backup server or storage node.
   - Configure your Solstice Backup server or storage node as an administrator, so you can enter the `dasadmin allocd` command to allocate devices from your Solstice Backup server or storage node computer.

▼ *To Find the Assigned Names for Devices in the Silo*

To find the names assigned to the devices in the silo, you can use a utility called `dasadmin` that is supplied with the Solstice Backup software.

To find the names:

1. Set three environment variables:
   - DAS_SERVER, the hostname of the silo management server, which runs DAS
   - DAS_CLIENT, the hostname of the Solstice Backup server or storage node
   - ACI_MEDIA_TYPE, one of the following values: 3480, CD_THICK, CD_THIN, DECDLT, 8MM, 4MM, D2, VHS, 3590, CD, TRAVAN, DTF, BETACAM, AUDIO_TAPE, or DAS_MEDIUM, same value as ACI_MEDIA_TYPE

2. Enter the following command to see a list of drives and the hostnames to which they are allocated:

   # **dasadmin ld**

▼ To Define a Silo Device

If you add a new device to the silo and you want the Solstice Backup server to recognize and use the device, you must first delete the existing Jukebox resource, and then reconfigure your silo's Jukebox resource to access the new device.

To define a silo device:

1. Log on as root to the Solstice Backup server that uses the silo.

2. In the Solstice Backup Administrator program, select Jukeboxes from the Media menu to open the Jukeboxes window.

3. Select the silo listing in the Jukeboxes attribute.

4. Click Delete.

   When prompted to confirm the deletion of the silo resource, click OK.

5. Click Apply.

6. Run the jbconfig program to reconfigure the silo. For more information, see "To Add Autochangers" on page 361. Use exactly the same name for the silo as before, and add the additional device.

If you define multiple devices in a silo for the Solstice Backup software to use, you do not need to specify them in any particular order in the Jukebox resource. The only restriction is that the list of device names in the Devices attribute must be in the correct order with respect to the list in the STL Device Names attribute.

▼ To Release a Silo Device

When a silo device is configured for use with a Solstice Backup server, you have the option of restricting silo access only to the Solstice Backup server. These restrictions allow increased availability to the silo for those with full access. However, these restrictions can be lifted by using the Release Device feature.

To release a silo device:

1. At the main Solstice Backup Administrator program window, select the appropriate autochanger in the Devices window.

2. Select Media>Release Shared Device.

3. Repeat these steps for each device you want to release.

Silo Device Cleaning

Do not enable automated cleaning for the silo in the Solstice Backup Administrator program. To find out how to clean devices in your silo, refer to the documentation for your silo management software. The automated device cleaning feature cannot be used in a silo because it depends on fixed slot numbers.

▼ To Set Environment Variables for EMASS/Grau and StorageTek Silos

For EMASS/Grau and StorageTek silos, you need to set the following environment variables in the Solstice Backup startup script after the lines:

```
(echo 'starting Solstice Backup daemons:') > /dev/console
```

Solstice Backup startup scripts are located in:

TABLE 10-4   Solstice Backup Startup Script Location

Operating System	Startup Script Location
Solaris, DYNIX/ptx, and IRIX	/etc/init.d/networker
HP-UX	/sbin/init.d/networker
Linux	/etc/rc.d/init.d/networker
AIX	/etc/rc.nsr

Enter the environment variables using the format:

```
ENV_VAR = value
export ENV_VAR
```

The following environment variables should be set:

TABLE 10-5 EMASS/Grau and StorageTek Silo Environment Variables

Silo Model	Environment Variables
Emass/Grau	DAS_SERVER = *name-of-DAS-server* DAS_CLIENT = *name-of-system-as-defined-to-DAS-server* DAS_MEDIUM = *type-of-tape-drive-used* ACI_MEDIA_TYPE = *type-of-tape-drive-used* For DAS_MEDIUM and ACI_MEDIA_TYPE, use one of the following values: • 3480 • VHS • OD-Thick • 3590 • OD-Thin • CD • DECDLT • TRAVAN • 8MM • DTF • 4MM • BETACAM • D2 • AUDIOTAPE
StorageTek	CSI_HOSTNAME = *name-of-ACSLS-system* Additionally, the following commands should be running on your system and can be included in the Solstice Backup startup script: /sbu_*binaries_path*/mini_el & /sbu_*binaries_path*/ssi &

Media Management in a Silo

More than one software application can use a single silo. Therefore, media management in a silo requires some extra operations to protect the volumes used by other programs from being overwritten by the Solstice Backup software.

This section addresses how volumes are specified for Solstice Backup software to use, how volumes are mounted in the devices, and how volumes are identified and counted in a silo.

Silo Slot Numbers

In an autochanger, Solstice Backup software specifies many of the functions by a slot number. A silo uses the same idea. In an autochanger, there is a fixed number of slots, and the Solstice Backup software uses the slot number to refer to the physical location of a volume.

A silo, however, has a variable number of slots, starting at zero when it is first configured and limited by the silo license you purchased. The fundamental identifier of a silo volume is its bar code, or *volser* (volume serial number). The volser never changes over the life of a particular volume.

When the nsrjb command displays a list of the contents of a silo, it also displays a slot number. You can use the slot number to specify which volumes to mount, unmount, label, and inventory. Volumes are not always assigned the same slot number in the silo. The slot numbers in the silo are assigned dynamically, based on the sorted order of the bar codes that have been allocated. If you allocate additional bar codes that fall earlier in the sort sequence, the slot numbers of all the volumes later in the sequence change. For more information, refer to the *Legato Command Reference Guide* or the nsrjb man page.

Because the slot number is not a perfect identifier for a silo volume, operations that might change the slot number cannot accept slot numbers as arguments. For example, you cannot deallocate volumes based on slot numbers, because this operation can change the slot numbers of volumes being deallocated.

Using Silos with a Volume Import/Export Capability

The Solstice Backup software supports the use of the SCSI-II import/export feature found in many brands of silos. Depending on the silo model, this feature is also known as cartridge access port (CAP), mail slot, and loading port. The import/export feature deposits and withdraws volumes from slots in the silo.

The import/export feature enables the operator to deposit and withdraw cartridges without invalidating the device inventory list. Normally, if the operator opens the door to load or unload volumes, the element status of the autoloader is invalidated, requiring the long operation of reinitializing the silo. Note, however, that Solstice Backup software does not automatically inventory the volume after a deposit.

You can use Solstice Backup commands or the silo management software to control the import/export feature on the currently supported silos to deposit and withdraw volumes in a silo. However, it is often more efficient to use the silo management software, especially to deposit or withdraw a large number of volumes.

Unlike autochangers, you cannot issue deposit and withdraw commands using the Solstice Backup Administrator program. Rather, you need to enter the appropriate commands at the command prompt:

■ The Solstice Backup command to allocate and deposit volumes is

```
# nsrjb -a -T tags -d
```

   where *tags* specifies the tags or bar codes of volumes in a remote jukebox.

■ The Solstice Backup command to deallocate and withdraw volumes is

```
# nsrjb -x -T tags -w
```

   where *tags* specifies the tags or bar codes of volumes in a remote jukebox.

On some silos (for example, StorageTek silos with the import/export feature set to automatic mode), the silo management software inserts volumes automatically. You cannot use Solstice Backup software to insert volumes.

On StorageTek silos, due to differences between the internal operations of the Solstice Backup software and the StorageTek silo management software, the Solstice Backup software can only withdraw one volume at a time. You must physically remove the volume from the silo's import/export area before any more volumes can be withdrawn.

Mounting and Unmounting a Silo Volume

Mounting and unmounting silo volumes is the same as for autochangers:

■ You must mount a volume before you can label it or read or write data on it. The robotic mechanism mounts volumes in the devices of a silo.

■ You must unmount volumes before you can inventory the volumes in a silo or deallocate a volume from a Solstice Backup pool.

For details on mounting and unmounting volumes, see "Remote Autochanger Management" on page 372.

To specify a bar code identifier or template for the volumes, use the -T option with the nsrjb command. For more information, refer to the *Legato Command Reference Guide* or the nsrjb man page.

Labeling a Silo Volume

The Solstice Backup labels for volumes in a silo include both a regular Solstice Backup volume label (written on the media of the volume) and a silo bar code identifier. The volume label is usually based on the volume pool's label template. The bar code identifier is written on a physical label on the outside of the volume, which the bar code reader in the silo can scan during inventory. For instructions on how to label silo volumes, see "To Label Autochanger Volumes" on page 395.

The use of bar codes, and matching bar code labels and Solstice Backup volume labels, are both enabled for a silo. You can change the Match Bar Code Labels attribute, but do not set the Bar Code Reader attribute to No. When both attributes are set to Yes, the internal volume label that Solstice Backup software writes on the media of each volume matches the bar code label on the outside of the volume. When the labels match, it is easier to track volumes, but the Solstice Backup software does not require the internal and external labels to match.

With most silo management software, you can use unlabeled volumes. The silo management software assigns a "virtual" bar code label to those volumes. Although you can use volumes without bar codes, it is very difficult to maintain integrity, since after you remove the volume from the silo, the information about the virtual bar code is lost. Any volume without a bar code can be reinserted into the silo under a virtual bar code that Solstice Backup software (or another application) associates with some of your data.

Allocating a Silo Volume

When you allocate volumes, you direct the Solstice Backup server to those volumes it can use.

Caution – Because silos can be used by more than one software application, it is possible that a different application could read or write to volumes that belong to the Solstice Backup software. To prevent this problem, most silo management software includes methods to limit access to volumes based on the host name of the computer on which Solstice Backup software and the other programs run. The Solstice Backup software does not provide any method for setting up this sort of protection; the silo management software must configure it.

Allocating a volume causes Solstice Backup software to query the silo management software to verify that the volume you requested exists.

- If the volume exists, the volume is allocated to the Solstice Backup software.
- If the volume does not exist, the following message appears:

```
barcode xyz is not present in the silo and was not added
```

If you are allocating a range of volumes, the allocation continues after displaying the message. The message is informational and does not indicate a Solstice Backup error.

▼ To Allocate a Silo Volume

To allocate silo volumes:

1. In the Solstice Backup Administrator program, select Jukeboxes from the Media menu to open the Jukeboxes window.

2. Select Media>Add Library Volumes to open the Add Library Volumes window.

3. Select the method (Template or List) of entering the names of the volumes to allocate, and then enter their names in the Barcodes attribute.

4. Click OK.

To deposit volumes into a silo and then allocate them (on silos that require manual depositing, like EMASS/Grau), place the volumes in the insert area, then enter the following command:

```
# nsrjb -a -T tags -d
```

where *tags* specifies the tags or bar codes of volumes in a remote jukebox.

On StorageTek silos, the silo management software deposits volumes automatically.

For more information on the silo tape library, see "How Solstice Backup Software Interacts with a Silo" on page 407.

Deallocating a Silo Volume

If you no longer need an STL volume in a silo for the Solstice Backup software, you can deallocate the volume from the library. Deallocation is basically the same operation as removing a volume from an autochanger. Although the volume cannot be loaded by the robotic mechanism, the entries in the Solstice Backup media database remain intact. If you allocate the volume again later, Solstice Backup software can retrieve the data from it.

Use deallocation when your silo license limits the number of slots you can use, or when you move data offsite for safer storage. If your license limits the number of slots, you might be able to leave the volumes in the silo. This way, you can easily reallocate the volumes when you need to access the data on them.

The allocation operation is not automatic. You must manually allocate the volumes again and reinventory them to let the Solstice Backup server access the data. If you remove the volume from the silo for offsite storage, you must deallocate it with the Solstice Backup software and then use the silo management software to eject it from the silo.

Caution – StorageTek silos currently can eject only one volume at a time. The silo operator must remove each volume before another `nsrjb -x -w` command can be issued. If you deallocate and withdraw multiple volumes, they are all deallocated, but only the first is ejected.

▼ To Deallocate a Silo Volume

To deallocate silo volumes:

1. Unmount the volume from the device. For instructions on unmounting volumes, see "To Unmount Volumes" on page 394.

2. In the Solstice Backup Administrator program, select Jukeboxes from the Media menu to open the Jukeboxes window.

3. Select Media>Remove Library Volumes to open the Remove Library Volumes window.

4. Select the method (Template or List) of entering the names of the volumes to deallocate, and then enter their names in the Barcodes attribute.

5. Click OK.

For more information on STL, see "How Solstice Backup Software Interacts with a Silo" on page 407.

NDMP Support for the Solstice Backup Server

This chapter provides instructions on how to enable and use the NetWorker NDMP Connection.

This chapter includes the following sections:

Introduction to NDMP

Network Data Management Protocol (NDMP) is a standards-based, storage management client/server protocol. It allows the Solstice Backup server to provide connections to computers with NDMP data modules for tape operations. It also allows the Solstice Backup server software to maintain the client files and media indexes, as well as to control backup operations on a computer with the NDMP service installed.

NDMP Connection Support

The NetWorker NDMP Connection allows the Solstice Backup server to communicate with other computers on the network that support NDMP. With the NDMP Connection, you can connect to computers that have an active NDMP service and an NDMP data module installed. The computers, except for the server, do not need to have the Solstice Backup software installed.

TABLE 11-1 lists the terms used with NDMP.

TABLE 11-1   NDMP Terms

Term	Definition
NDMP	A client/server protocol that allows any storage management application, such as Solstice Backup software, to control backup operations on a computer.
NDMP service	A network communication service that uses NDMP. This service can include a variety of additional features, such as backup and retrieval utilities.
NDMP data module	An agent used to drive an NDMP service. The Legato Celestra product is an example of an NDMP data module or agent.
Application server	The computer where the NDMP service and the NDMP data module reside. This computer controls the data to be backed up.
NDMP tape server	The application server with the tape device attached.
NDMP computer	Any computer with an NDMP data module or agent installed.

The NDMP Connection, part of the Solstice Backup server software, is distributed on CD-ROM.

To install an NDMP data module or enable NDMP support on an application server, refer to the respective vendor's documentation.

Licensing the NetWorker NDMP Connection

NetWorker NDMP Connection is an optional product that must be licensed separately from the Solstice Backup base product. An NDMP client connection must be applied to each NDMP client. For more information, the Solstice Backup Roadmap.

NDMP Support for the Solstice Backup Server

The following sections describe requirements and considerations for using NDMP services.

NDMP Connection Software and Installation Requirements

Before you enable the NDMP Connection, you must install Solstice Backup server software on a Solaris computer designated for this purpose, and you must be using Solstice Backup Network Edition or Solstice Backup Power Edition.

For more information on installing the Solstice Backup software, refer to the *Sun Solstice Backup Installation Guide*.

The NetWorker NDMP Connection is supported on Solaris versions 2.6, 2.7 (32- and 64-bit) (UFS/VxFS).

In addition, the Legato SnapImage Module software is supported to allow remote backups. For information about the SnapImage Module software, refer to the *Legato SnapImage Module Installation and Administrator's Guide*. For information about configuring a server or storage node with SnapImage Module software for remote backups, see to .

NDMP Client Support

To request backup and recovery services from the Solstice Backup server, Solstice Backup clients must be able to access the Solstice Backup software.

You do not need to install Solstice Backup client or storage node software if you plan to back up application servers through NDMP only. If the application server is also a Solstice Backup server, the Solstice Backup client and storage node software is installed as part of the server installation.

If you plan to back up the computers with standard Solstice Backup backups in addition to the NDMP backups, you must install Solstice Backup client software on each computer and Solstice Backup storage node software on each computer with a tape device.

Silo Support

Silos controlled by either ACSLS or DAS controllers with the NDMP-enabled Solstice Backup server are supported. Any silo behind the ACSLS or DAS is supported, since all robotic control and communication is done by ACSLS and DAS, which receive commands through TCP/IP from the Solstice Backup server.

Snapshot Backup Support

NetApp creates two types of filesystem snapshots—those taken at specified intervals and those taken when a backup occurs. The snapshots taken during a backup are automatically deleted when the backup is completed.

Solstice Backup 6.0 and later software supports backing up and recovering the snapshots that are taken at the specified intervals and which are stored in the snapshot directory. For more information about snapshots, refer to the NetApp documentation.

Limitations of NDMP

The Solstice Backup software has the following limitations when using NDMP:

■ To back up the bootstrap file, you must have a non-NDMP device attached to the Solstice Backup server. The bootstrap includes the media index and configuration files needed for recovering the Solstice Backup server after a disk crash.

■ File type devices are not supported as NDMP devices.

■ In order to perform either a backup or recovery operation, you must be logged in as root.

■ Incremental backups are not supported. If you select the incremental setting, the Solstice Backup server performs a full backup. However, you can schedule differential level backups to function like incremental backups. For example, a weekly backup schedule of full on day 1 and incrementals on days 2 through 7 is the same as a weekly backup schedule of full on day 1, level 1 on day 2, level 2 on day 3, level 3 on day 4, and so on.

Note: If you are performing a Celestra backup with NDMP, only full and level 1 backups are supported. For more information about Celestra backups, refer to the Celestra documentation.

Configuration Tips

The following sections provide configuration tips and workarounds to ensure a successful NDMP backup or recovery.

NDMP User Name and Password Consideration

If you use more than one user name, more than one password, or a blank for any client instance, you might experience recovery errors. If this happens, modify the NDMP user name and password fields for all clients to use a single user name and a single password.

Important Notes on Swap Space

If you want to back up or retrieve a large number of files, you may need to increase the amount of swap space. Backups and retrievals fail if there is an inadequate amount of swap space.

- For backups, the Solstice Backup server uses 100 MB of swap space for every one million files.
- For retrievals, the application server uses 350 MB of swap space for every one million files.

If the backup is still running slowly after increasing the swap space, try increasing the amount of memory on the system.

Adjusting NDMP Load Sleep Value

When using an autochanger, it may take longer for the drive to become ready after loading a tape into that drive. When this occurs, you may receive one of the following error messages:

```
failed to establish connection to NDMP server
nsrmmd: failed to open ndmp tape device

no tape in specified device
```

If you receive one of these messages, adjust the Load Sleep attribute by completing the following steps:

1. From the nwadmin GUI, select Media>Jukeboxes to open the Jukeboxes window.

2. Select View>Details to display the hidden attributes.

3. For the Load Sleep attribute, increase the value to 40 or 45 seconds.

Note – The Load Sleep attribute may need to be increased more, depending on your configuration. If 40 or 45 seconds is not enough time, increase the value until you no longer receive the error message.

Modifying the maxdsiz Parameter in HP-UX to Accommodate Large Backups

On HP-UX systems, verify that the Process Maximum Heap Limit parameter for the kernel (maxdsiz) is large enough to accommodate backups of one million files or more. You do this by changing the maxdsiz parameter to 2 GB using the `sam` utility.

To change the maxdsiz parameter:

1. Log in as *root and* run the `sam` utility.

2. Select the Kernel Configuration option.

3. Select Configurable Parameters.

4. Select the maxdsiz parameter.

5. Change the value of the maxdsiz parameter to 2 GB.

6. Recompile and reboot the kernel.

Optimizing Qtree Backups with NetApp Data ONTAP

If you are using Data ONTAP 5.3.*x* and are backing up more than eight qtrees on one volume, you might want to use the NDMP **logical find** option to speed up the backup.

Caution – If you are running Data ONTAP 6.0 or later, do not use the **logical find** option. The filer automatically selects the optimal backup method.

To enable the **logical find** option, enter the following command at the filer console's prompt:

```
filer> ndmpd logical_find
```

When the option is enabled, the following message appears:

```
logical find set to: true
```

If the backup speed is not increased by enabling the **logical find** option, disable the option. To disable the option, enter the following command at the filer console's prompt:

```
filer> ndmpd logical_find
```

When the option is disabled, the following message appears:

```
logical find set to: false
```

Note – If you reboot the filer, you must enable the **logical find** option again.

For more information regarding qtrees, refer to the Data ONTAP documentation.

Backing Up Large Files with Auspex

By default, support for files greater than 2 GB is not enabled for an Auspex backup, allowing usage of a standard gtar tape format. To enable support for large files:

1. Open the /usr/AXndmp/etc/config file.

2. Remove the comment marks for the line containing ALLOW_TAR_EXTENSION.

3. Save the modified file.

You can now back up large files.

Note – If you modify your system to support large files, the tapes will be in nonstandard gtar format.

Setup and Configuration for NDMP Operations

This section assumes that you have completed installation of your Solstice Backup software.

After installation, you need to perform these tasks:

- Enable and register the NDMP Connection.
- Install and configure an NDMP data module or NDMP-enabled application server to work with the Solstice Backup server.
- Configure your autochanger or tape device.
- Configure the Solstice Backup server for NDMP operations.
- Run a test backup and retrieval to be sure your configuration is correct.

contains a setup and configuration roadmap.
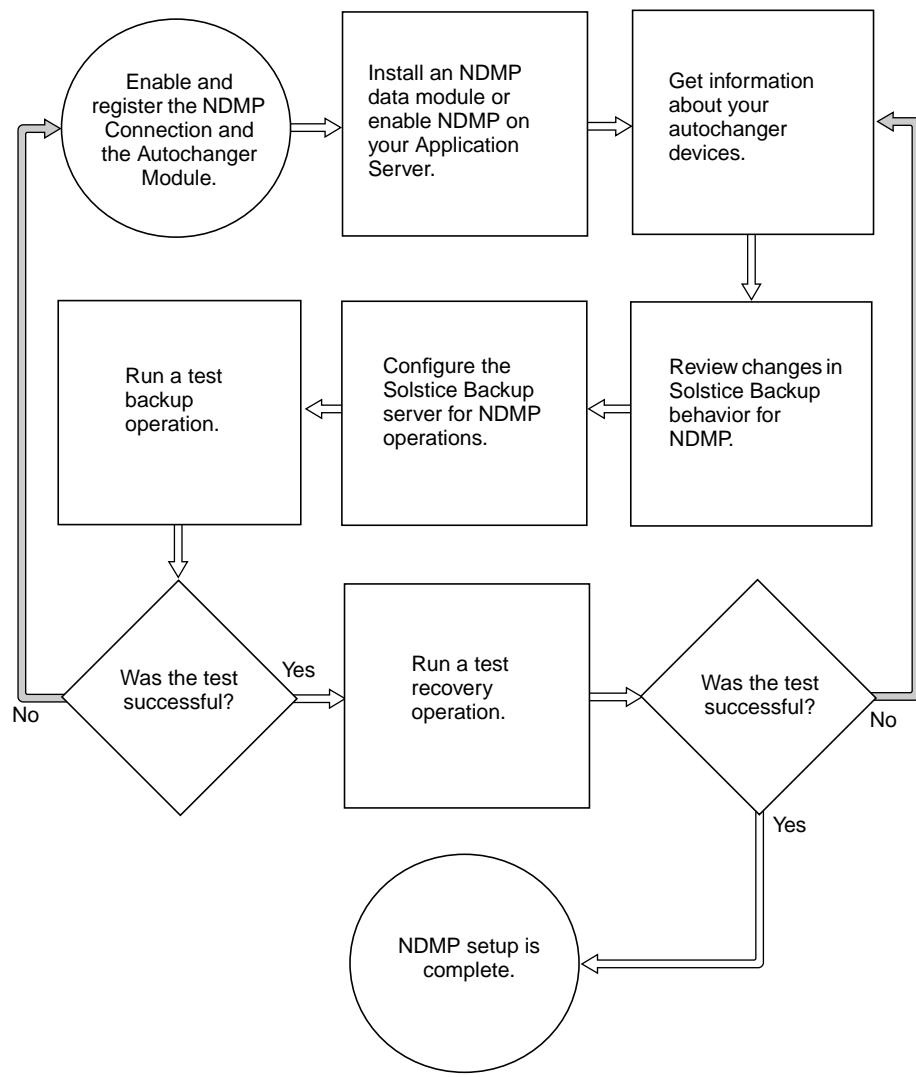
FIGURE 11-1 Setup and Configuration Roadmap

Configuring Autochanger Support for NDMP Operations

In order to perform NDMP operations, you must configure your autochangers by running the **jbconfig** command. To do this, you must first obtain the NDMP autochanger handle and the pathname of the media device. You obtain the required information by issuing the appropriate commands for your tape server.

You can also create a file for NDMP devices, called *ndmpjbconf_NDMPhostname*, that can later be searched using the Solstice Backup **inquire** program. If you do not create the file, you cannot use the **inquire** program for NDMP devices. However, you do *not* need to create the file to set up your autochangers or run the **jbconfig** program. The **inquire** program is optional for users who prefer to use it.

Obtaining Information about NDMP Autochangers

Before you can configure your autochanger for NDMP operations, you must tell the Solstice Backup server the *NDMP autochanger handle* and the pathname of the media device.

Before you set up your autochanger or configure the Solstice Backup server for NDMP operations, enable NDMP support on your application server.

- For SnapImage Module devices, see "To Find the Autochanger Handle and Device Pathname for a SnapImage Module Device" on page 429.
- For NetApp devices, see "To Find the Autochanger Handle and Device Pathname for a NetApp Computer" on page 431.
- For EMC devices, see "To Find the Autochanger Handle and Device Pathname for an EMC Computer" on page 433.
- For Auspex devices, see "To Find the Autochanger Handle and Device Pathname for an Auspex Computer" on page 434.

▼ To Find the Autochanger Handle and Device Pathname for a SnapImage Module Device

The required steps for determining the autochanger handle and device pathname are different between Solaris and HP-UX computers. Complete the instructions for your appropriate platform.

▼ *To Find the Autochanger Handle and Device Pathname for a Solaris Computer*

To identify the autochanger handle and the device pathname for a Solaris computer, use the `ls -l` command to match their SCSI IDs.

1. To determine the autochanger handle, enter the following command at the prompt:

```
# ls -l /dev/rsjb*
```

The computer responds with a list of *rsjb* device handles:

```
lrwxrwxrwx  1 root     other      45 Jan 11 13:35 /dev/rsjb2 -
> ../devices/pci@1f.0/scsi@2,1/sjb@1,0:8r
```

- The autochanger robotics is identified by the `:8r` at the end.
- The SCSI ID field for the rsjb device in this example is `/sjb@1, 0`.
- The autochanger handle is `/dev/rsjb2`.

2. To determine the media device pathname, enter the following command at the prompt:

```
# ls -l /dev/rmt/*
```

The computer responds with the list of device filenames:

```
lrwxrwxrwx  1 root     root       87 Nov 17 12:03 /dev/rmt/0mbn
-> ../../devices/pci@1f.0/scsi@2/st@1,0:mbn
```

- The SCSI ID field for the media device in this example is `/st@1, 0`.
- The media device pathname is `/dev/rmt/0mbn`.

With the autochanger handle and media device pathname, you can configure your NDMP autochanger and create the *ndmpjbconf_NDMPhostname* file.

■ To create the *ndmpjbconf_NDMPhostname* file, which allows you to use the `inquire` utility to find NDMP devices, continue with .

You do *not* need to create the file to initially set up your autochangers or run the **jbconfig** command.

■ To configure the NDMP autochanger before you create the *ndmpjbconf_NDMPhostname* file, skip to .

▼ *To Find the Autochanger Handle and Device Pathname for an HP-UX Computer*

To identify the autochanger handle and the device pathname for an HP-UX computer, use the ioscan command or the sam utility.

■ To use the ioscan command, at the prompt, enter the following command:

```
# ioscan -fnC tape
```

The autochanger handles and device pathnames are listed.

■ To use the sam utility:

a. Start the utility.

b. In the GUI, select Peripheral Devices.

c. In the next window, select Tape Drives.

d. Highlight the appropriate tape drive and select Action>Show Device Files.

The required information is displayed.

With the autochanger handle and media device pathname, you can configure your NDMP autochanger and create the *ndmpjbconf_NDMPhostname* file.

■ To create the *ndmpjbconf_NDMPhostname* file, which allows you to use the inquire utility to find NDMP devices, continue with . You do *not* need to create the file to initially set up your autochangers or run the **jbconfig** command.

■ To configure the NDMP autochanger before you create the *ndmpjbconf_NDMPhostname* file, skip to .

▼ To Find the Autochanger Handle and Device Pathname for a NetApp Computer

Use the sysconfig command to identify the autochanger handle and device pathname.

Caution – The procedure that follows is only an example. Check your NetApp documentation for the correct syntax and usage.

To identify the autochanger handle and media device pathname for a NetApp computer:

1. To verify that the tape device is installed on the NetApp filer, enter the following command at the prompt:

```
# sysconfig -v
```

The computer responds with:

```
NetApp Release 5.3.5: Sun Mar 12 13:29:37 PST 2000

        System ID: 0016784666 (NetAppHost1)

                slot 6: SCSI Host Adapter 6 (Qlogic ISP 1040B)

                Fast/Wide, Differential

                Firmware Version 4.65 Clock Rate 60MHz.

                4: QUANTUM Powerstor L200   001F

                5: QUANTUM DLT7000          2255
```

2. To determine the autochanger handle, enter the following command at the prompt:

```
# sysconfig -m
```

The computer responds with the devices on your computer:

```
Media changer (6.4)  QUANTUM  Powerstor L200

mc0 - media changer device
```

The autochanger handle is *mc0*.

3. To determine the device pathname, enter the following command at the prompt:

```
# sysconfig -t
```

The computer responds with a list of media device names.

```
Tape drive (6.5)  Digital DLT7000

nrst0a -  no rewind device, format is: 85937 bpi 70 GB (w/comp)

nrst0m -  no rewind device, format is: 85937 bpi 35 GB
```

- Where (6.5) indicates slot 6 and SCSI ID 5. This information must match the output in the **sysconfig -v** command.
- The media device name is nrst0a.

With the autochanger handle and media device pathname, you can configure your NDMP autochanger and create the *ndmpjbconf_NDMPhostname* file.

- To create the *ndmpjbconf_NDMPhostname* file, which allows you to use the inquire utility later to find NDMP devices, continue with . You do *not* need to create the file to initially set up your autochangers or run the **jbconfig** command.

- To configure the NDMP autochanger before you create the *ndmpjbconf_NDMPhostname* file, skip to .

▼ To Find the Autochanger Handle and Device Pathname for an EMC Computer

Use the **list** command to identify the autochanger handle and device pathname.

Caution – The procedure that follows is only an example. Check your EMC documentation for the correct syntax and usage.

To identify the autochanger handle and device pathname, enter the following on the command line:

Celerra> **server_devconfig** *hostname* **–list –scsi –nondisks**

You may also obtain the autochanger handle through the browser GUI.

To access the Administrator GUI:

1. In the browser window, enter:

   http://EMCcontrolhostIP:8000/top_level.htm.

2. At the prompt, enter your user name and password.

3. Use this EMC admin GUI to list and/or determine SCSI devices.

With the autochanger handle and media device pathname, you can configure your NDMP autochanger and create the *ndmpjbconf_NDMPhostname* file.

- To create the *ndmpjbconf_NDMPhostname* file, which allows you to use the inquire utility later to find NDMP devices, continue with "To Create the ndmpjbconf_NDMPhostname File" on page 435. You do *not* need to create the file to initially set up your autochangers or run the **jbconfig** command.

- To configure the NDMP autochanger before you create the *ndmpjbconf_NDMPhostname* file, skip to "To Configure an NDMP Autochanger" on page 437.

▼ To Find the Autochanger Handle and Device Pathname for an Auspex Computer

The autochanger handle and device pathname for an Auspex computer can be found in the message file. To locate this information, open the /var/adm/messages file.

Note: For information about adding a SCSI device, refer to your Auspex documentation.

The file contains the list of devices, for example:

```
Jan 18 10:32:22 nightcap FSP_0[1611]:fsp0c7t0:STK      9714      [
Medium_Changer ][ Fast ] [ min:1792 ]

Jan 18 10:32:22 nightcap FSP_0[1611]:fsp0c7t1:Quantum DLT4000  [
Tape ][ Not Negotiated ] [ min:1808 ]

Jan 18 10:32:22 nightcap FSP_0[1611]:fsp0c7t2:Quantum DLT4000  [
Tape ][ Not Negotiated ] [ min:1824 ]
```

In this example, the following information is obtained:

handler: **fsp0**

pathname drive 1: **/dev/raxmt/fsp0c7t1n**

pathname drive 2: **/dev/raxmt/fsp0c7t2n**

With the autochanger handle and media device pathname, you can configure your NDMP autochanger. For instructions, see "To Configure an NDMP Autochanger" on page 437.

▼ To Create the ndmpjbconf_NDMPhostname File

The *ndmpjbconf_NDMPhostname* file allows you to use the inquire command to obtain information about the NDMP autochanger. The file is not required to initially set up your NDMP autochanger.

The *ndmpjbconf_NDMPhostname* file is a one-line text file that contains:

- NDMP hostname
- User name
- NDMP autochanger handle

The format of the file is:

```
server user autochanger_handle
```

At the end of the configuration, you are notified of the location of the created file.

The user name must match the client and device/autochanger resources that you will set up later as you configure the Solstice Backup server for NDMP backups. The user name is the NDMP user account name for the computer where the autochanger robotics are installed.

To create the *ndmpjbconf_NDMPhostname* file:

1. Make sure you have the autochanger handle and device pathname information for the NDMP autochanger. To get more information, see "Obtaining Information about NDMP Autochangers" on page 429.

2. Log in as administrator on the Solstice Backup server.

3. Enter the following command at the prompt:

```
# ndmpjbconf
```

The script uses the term *jukebox* instead of *autochanger*. The terms are interchangeable.

4. Provide the information as requested:

```
Enter NDMP Server name: host1[Return]

Enter NDMP user name: root [Return]

Enter NDMP jukebox handle: /dev/nrst0a[Return]
```

The configuration files are created in /usr/lib/nsr. You are notified of the location when ndmpjbconf is finished.

▼ To Use the inquire Command with NDMP Devices

Once you have created the *ndmpjbconf_NDMPhostname* file, you can use the inquire command to search for available devices on the NDMP server. New options have been added to inquire to support NDMP devices.

Note – The inquire command can retrieve only information in the *ndmpjbconf_NDMPhostname* file. It cannot retrieve device filenames for NDMP devices.

To get NDMP autochanger information from the location on the local server:

1. Enter the following command at the prompt:

   # **inquire -N** *NDMP_type*

   Where *NDMP_type* is either *standard*, *celestra*, *celestra-hp*, or *netapp*.

Note – If you are configuring an EMC or Auspex device, the *NDMP_type* is *standard*.

You are prompted with:

```
# Enter 'root' password on NDMP host 'hostname' (characters
will not be echoed)
```

2. Enter the required password.

The inquire command shows devices as defined in the *ndmpjbconf_NDMPhostname* file.

To get NDMP autochanger information on a specific NDMP host when the file is not in the default location:

1. Enter the following command at the prompt:

   # **inquire -N** *NDMP_type* **-f** *absolute_path_of_ndmpjbconf_NDMPhostname*

   The -f option is only valid with the –N option. It cannot be used alone.

   You will be prompted with:

   # Enter 'root' password on NDMP host '*hostname*' (characters will not be echoed)

2. Enter the password.

The inquire command shows devices on the specified NDMP host.

Note – If neither NDMP_type nor the absolute path of the *ndmpjbconf_NDMPhostname* file is provided, the inquire command shows devices on the local Solstice Backup server.

For more information on the inquire command, refer to the inquire man page.

▼ To Configure an NDMP Autochanger

The Solstice Backup autochanger resource configuration tool, jbconfig, helps you configure your autochangers for NDMP operations.

To configure your autochanger:

1. Make sure you have the autochanger handle and device pathname information for the NDMP autochanger. To get this information, see "Obtaining Information about NDMP Autochangers" on page 429.

2. Log in as root on the Solstice Backup server.

3. Enter the following command at the prompt:

   ```
   # jbconfig
   ```

   Note: You must know the NDMP password to use **jbconfig**.

4. Select *Install an Autodetected NDMP SCSI Autochanger* or *Install an Autodetected SCSI Autochanger* when prompted and respond to the questions.

   To determine which option to select, consider the following:

- If the autochanger is connected to an NDMP tape server and the robot is controlled by the NDMP tape server, select the Autodetected NDMP SCSI Autochanger option.

- If the Solstice Backup backup server or storage node controls the robot and shares the autochanger with an NDMP tape server, select the Autodetected SCSI Autochanger option. Configure the tape drives to be used by the NDMP tape server.

An example that shows how to configure an NDMP autochanger for a NetApp filer follows.

Note – The script uses the term *jukebox* instead of *autochanger*. The terms are interchangeable.

Note – This example might not be valid for your NDMP server. Refer to the documentation that came with your NDMP server to identify valid device names.

Example: How to Configure an NDMP Autochanger for a NetApp Computer

```
# jbconfig

1) Install a SmartMedia Jukebox.

2) Install an Autodetected SCSI Jukebox.

3) Install an Autodetected NDMP SCSI Jukebox

4) Install an SJI Jukebox.

5) Install an STL Silo.

What kind of Jukebox are you installing? [1] 3 [Return]

Enter NDMP Server name: Host1 [Return]

Enter NDMP user name: root [Return]

Enter NDMP user password (characters will not be echoed):

Enter NDMP Jukebox handle:? mc0[Return]

What is the NDMP type of 'host1'?

1) One of the Standard NDMP Servers.

2) NetApp or Celestra HP.

Choice? 2 [Return]

Communicating to devices on NDMP Server 'host1', this may take a
while...

These are the SCSI Jukeboxes currently attached to your system:

1) scsidev@1027.0.0: Standard SCSI Jukebox, Vendor,<ATL>, Product
<L500 6320000>:

Which one do you want to install? 1 [Return]
```

```
Installing a 'Standard SCSI Jukebox' jukebox

Name you would like to assign to the jukebox device? *atl* [Return]

A brief description of the jukebox device? [scsidev@1027.0.0:
Standard SCSI Jukebox [Return]

Pathname of the control port for the jukebox device?
[scsidev@1027.0.0] [Return]

Do you want automated device cleaning support enabled? [yes/no)
**yes** [Return]

Enter hostname that owns media drive 1: ? *host1* [Return]

An example of a valid device name for this OS is 'nrst0a'

Will media drive 1 be shared by multiple device paths? (yes/no)
**no**[Return]

Enter pathname of media drive 1:? *nrst0a* [Return]

using 'rd=host1:nrst0a' as device name

Should the drive be configured as an NDMP device? (yes/no) **yes**
[Return]

This media device has not been configured yet. Please

select a media device type for rd=host1:/dev/rmt/0mbn.

        1) 3480                 22) dst

        2) 3570                 23) dst (NT)

        3) 3590                 24) dtf

        4) 4890                 25) file

        5) 4mm                  26) himt

        6) 4mm 12GB             27) logical
```

```
        7)  4mm 20GB                    28)  LTO Ultrium

        8)  4mm 4GB                     29)  optical

        9)  4mm 8GB                     30)  qic

       10)  8mm                         31)  SD3

       11)  8mm 20GB                    32)  sdlt

       12)  8mm 5GB                     33)  SLR

       13)  8mm AIT                     34)  tkz90

       14)  8mm AIT-2                   35)  tz85

       15)  8mm Mammoth-2               36)  tz86

       16)  9490                        37)  tz87

       17)  9840                        38)  tz88

       18)  dlt                         39)  tz89

       19)  dlt1                        40)  tzs20

       20)  dlt7000                     41)  vhs

       21)  dlt8000                     42)  VXA

Choice? 20


Your jukebox does not report that it has a bar code reader, but it
does report that it can handle volume tags. Some jukeboxes that
have barcode readers report this way. If your jukebox has a barcode
reader, do you want bar code reader support enabled? (yes/no) no

Jukebox has been added successfully

Would you like to configure another jukebox? (yes/no) no
```

If you are sharing a drive using the DDS feature, the **jbconfig** utility asks you for additional information about that drive. For an example, see "To Enable DDS When Configuring an Autochanger" on page 374.

For more information on the **jbconfig** and **nsr_jukebox(5)** commands, refer to the *Legato Command Reference Guide* or the **jbconfig** and **nsr_jukebox** man pages.

Sharing Autochangers between NDMP Clients

The Solstice Backup software permits the control and usage of individual devices within an autochanger to different Solstice Backup hosts within a datazone. For additional information about sharing autochangers, see "Sharing Autochangers between Solstice Backup Hosts" on page 379.

The Solstice Backup 6.1 software also includes support for DDS, which allows sharing autochangers in a SAN environment. You must, however, ensure that the NDMP devices are homogeneous, because DDS cannot be enabled between different types of NDMP devices.

As an example, DDS can be enabled in the following configurations:

- EMC Celerra to EMC Celerra
- NetApp to NetApp (any NetApp filers that Network Appliance supports within a zone)

For additional information about DDS, see "Dynamic Drive Sharing" on page 225.

How NDMP Affects Solstice Backup Software Functionality

Adding NDMP to your Solstice Backup datazone environment renders the following changes:

- NDMP operations require two new utilities and changes to a third:
  - nsrndmp_save, which controls backups on an NDMP computer, supports most of the same command-line arguments as save(1m). Additional options have been added for NDMP support.
  - nsrndmp_recover, which controls recover operations on an NDMP computer, supports most of the same command-line operations as recover(1m). Additional options have been added for NDMP support.

- nsrmmd, which adds the -r *hostname* option, indicates the hostname of the NDMP computer.

- You can browse other NDMP computers for recoveries through both the command line and the Solstice Backup Administrator program. You cannot browse NDMP computers for backups.

- Savegroup probes are not supported in the NDMP client resource. You must manually list the filesystems to be backed up with NDMP. You cannot use the default option, *All*, to choose filesystems to be backed up.

- To perform parallel backups, you must complete the following:

  a. Define one Group for each Client resource.

  b. Verify there is one tape device available for each Group.

  c. Schedule these Groups to start at the same time.

  d. Define separate pools for the NDMP backups and the non-NDMP backups.

  e. Verify there is a backup device on the Solstice Backup server so the indexes can be backed up locally.

- Cloning is not supported.

- Archiving is not supported.

- Multiplexing is not supported.

- Save set consolidation is not supported.

- The jbexercise utility is not supported. To verify that you can find the autochanger on the NDMP server, you can either reset the autochanger with the nsrjb -HE command or inventory the autochanger with the nsrjb -I command.

- The scanner command generates media database entries only for data backed up with NDMP. You cannot recover data with the scanner utility or regenerate the file index information. Instead, you must use special tape readers provided by your NDMP vendor.

- NSR resource attributes (nsrmmd **polling interval,** nsrmmd **restart interval, and** nsrmmd **control timeout)** do not apply to NDMP computers.

- The Save Set Recover options *Rename Recovered File* and *Discard Recovered File* are not supported by NDMP. *Current files are always overwritten*.

Configuring the Solstice Backup Server for NDMP Operations

Before you can back up data with NDMP, you must configure the Solstice Backup server for NDMP operations. Use the Solstice Backup Administrator program to configure the following resources on the Solstice Backup server before you back up data:

- Group resource

  Add the NDMP computer to a backup group. If you are performing parallel backups, you must define one group for each client.

- Schedule resource

  Select a schedule for the NDMP client.

- Policy resource

  Specify the browse and retention policies for the NDMP client.

- Pool resource and a Label Template resource

  You must specify separate pools for NDMP and non-NDMP volumes. Additionally, create a label template for NDMP operations.

- Client resource

  Create a Client resource with specific NDMP attributes. For instructions, see "To Configure an NDMP Client" on page 444.

- Device resource

  Create a Device resource for NDMP backups. For instructions, see "To Configure an NDMP Device Resource" on page 448.

When you configure the Solstice Backup server for NDMP operations, note that certain attributes of the resources are configured differently than for non-NDMP computers.

▼ To Configure an NDMP Client

For information on creating a Client resource, see "To Create a Client" on page 141. The following section covers issues relating to specific attributes of the Client resource that are affected by NDMP.

To create an NDMP client in the Client resource:

1. For the Name attribute, enter the client name.

2. For the Save Set attribute, list the filesystems to be backed up. You must specify the filesystem mount points or the raw partition names here. Savegroup probes and the default All option are not supported.

3. For the NDMP attribute, select Yes.

4. For the Remote Access attribute, list all computers that will store and retrieve entries for the client. In addition, list any computer that is used to browse entries for the NDMP computer. Since the Solstice Backup server stores the index entries for the NDMP computer, be sure that root@server is listed.

5. For the Remote User attribute, enter the NDMP account user name for the computer where the NDMP data module is installed. If this is root, be sure that root has a password.

6. For the Password attribute, enter the password for the NDMP user account. This password should match the password for the NDMP user account for the computer where the NDMP data module is installed.

7. For the Backup Command attribute, set the value to *nsrndmp_save*.

   The syntax for nsrndmp_save is:

   # **nsrndmp_save -T** *backup_type*

   The value for the NDMP backup type depends on the type of backup:

   - NetApp backups — enter **nsrndmp_save -T dump**
   - EMC backups — enter **nsrndmp_save -T pax**
     or **nsrndmp_save -T dump**
   - Auspex backups — enter **nsrndmp_save -T gtar**
     or **nsrndmp_save -T dd**

8. For the Application Information attribute, list the NDMP environmental variables. For a list of the variables and rules for entering their values, see "Specifying Application Information" on page 446.

9. For the Storage Node attribute, specify the appropriate storage node, depending on the type of backup:

   - For a local backup, enter the NDMP server name to which the tape device is attached. List this server first.
   - For a three-party or remote backup, list the destination server for the data being backed up first. The client backs up the data to each NDMP server or storage node in the order in which they appear in the attribute. For additional information about three-party backups, see "Three-Party Backup" on page 450. For additional information about remote backups, see "Remote Backup" on page 452.

10. Complete any other attributes as you would for a regular Solstice Backup client configuration.

Specifying Application Information

List environment variables that the Solstice Backup server uses for an NDMP backup. For the backups to be successful:

- Separate the environment variable name from its value with an equal (=) sign. For example:

  Application Information:
  HIST=*y*

- Put each entry in the Application Information attribute on a separate line.

- Click *Add* after each entry.

The values you enter depend on which NDMP data module or NDMP-enabled application server you use.

TABLE 11-2 lists required values for currently supported NDMP data modules and NDMP-enabled application servers. These values were correct at the time of publication. Check with the vendor for the latest information.

TABLE 11-2   Required Application Information Values *(1 of 2)*

NDMP Type	Required Values*	Means
Auspex	DIRECT=y/n	y/n. This value must be set to y when performing a direct access restore (DAR). If you are not performing a DAR recovery, the value is set to n. For additional information about DAR, see "Direct Access Restores" on page 461.
	HIST=y	y/n. Activates or deactivates file-by-file cataloging for gtar type. Default is n. If DIRECT=y, you must specify HIST=y.
	SNAPSHOT=y	y/n. Enables or disables snapshots. The default value is n.
	UPDATE=y	This value must be set to y. Forces /etc/dumpdates to be updated with the time, date, level, and filesystem being backed up. If a filesystem has previously been backed up at the same level as a new backup of the same filesystem, the old /etc/dumpdates entry for that filesystem will be overwritten with the new information.
EMC	HIST=y	y/n. Determines whether file history is created, which is required for index recovery. Default is n.
	UPDATE=y	This value must be set to y. Forces *etc/BackupDates* to be updated with the time, date, level, and filesystem being backed up. If a filesystem has previously been backed up at the same level as a new backup of the same filesystem, the old *etc/BackupDates* entry for that filesystem will be overwritten with the new information.
	emcfiles=	The files that will be backed up. Default is to leave the value blank.

TABLE 11-2 Required Application Information Values *(2 of 2)*

NDMP Type	Required Values*	Means
NetApp	DIRECT=y/n	y/n. This value must be set to y when performing a direct access restore (DAR). If you are not performing a DAR recovery, the value is set to n. DAR is supported on Data ONTAP 6.1 and later. For additional information about DAR, see "Direct Access Restores" on page 461.
	HIST=y	y/n. Determines whether file history is created, which is required for index recovery. Default is n. If DIRECT=y, you must specify HIST=y.
	UPDATE=y	This value must be set to y. Forces */etc/dumpdates* to be updated with the time, date, level, and filesystem being backed up. If a filesystem has previously been backed up at the same level as a new backup of the same filesystem, the old */etc/dumpdates* entry for that filesystem will be overwritten with the new information.

*These values were required by the NDMP data module or the NDMP-enabled application server at the time of publication. Check with the vendor 's web site for the most up-to-date values.

▼ To Configure an NDMP Device Resource

For details on how to configure the Device resource in the Solstice Backup software, see "Devices" on page 215. The following section provides instructions for configuring the Device resource information when using NDMP.

Note: The **jbconfig** utility creates the NDMP device for NDMP devices in a jukebox.

To configure a device:

1. For the Name attribute, specify the NDMP device using the following syntax:

```
rd=NDMP_computer:tape_device_pathname
```

2. Set the Target Sessions attribute to 1. Multiplexing is not supported on NDMP computers.

3. For the NDMP attribute, select Yes.

> **Note –** This option can only be set when you create the device. You cannot change the NDMP attribute after the device has been created. To change the device configuration, you must first delete the configuration, then recreate it.

4. For the Remote User attribute, enter the NDMP account user name for the computer where the NDMP autochanger robotics or tape device resides.

5. For the Password attribute, enter the password for the NDMP user account. This password should match the password for the NDMP user account for the computer where the NDMP autochanger robotics or tape device is installed.

6. Complete any other attributes as you would for a regular Solstice Backup device configuration.

Supported NDMP Backup Types

The workflow for NDMP backup and recovery operations is similar to the workflow for traditional Solstice Backup backup and recovery operations. Several utilities and several options have been added to the current utilities to enable NDMP operations.

The Solstice Backup server with NDMP Connection supports three configurations:

- Local — the backup is done from a tape device attached to the same server where the data resides (FIGURE 11-2 on page 450).
- Three-party — the backup is done from one NDMP server to another NDMP server (FIGURE 11-3 on page 451).
- Remote — the backup is done from an NDMP disk server to a Solstice Backup server or storage node with the NetWorker SnapImage Module software installed and tape drives attached (FIGURE 11-4 on page 453).

Note: If you run **nsrjb** to load a tape, you need to include the **-f** option and **rd** syntax in your device name. For example, suppose your filer name is netapp1 and your device is */dev/nrst0a*, and you want to load a tape from slot 1. The syntax for loading a tape into that drive is **nsrjb -l -f rd=netapp1:/dev/nst0a -S 1.**

Local Backup

The local backup involves a tape device attached directly to the server where the data resides. The data being backed up is sent directly from the NDMP server to its attached autochanger. Configure the Solstice Backup server for NDMP operations as you would for any NDMP backup.

FIGURE 11-2 illustrates a local configuration.

Solstice Backup Server   **NDMP Tape Server**   **Autochanger**

Legend			
Network Boundary	NDMP Connection	Meta Data	Backup Data
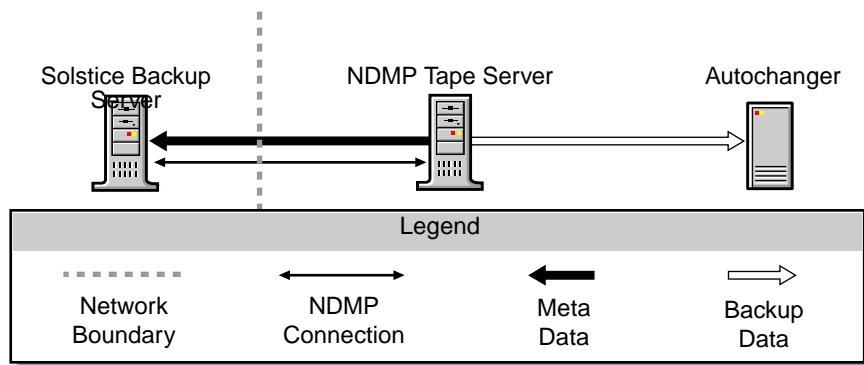
FIGURE 11-2 NDMP Local Configuration

Three-Party Backup

A three-party backup, also known as a three-way backup, is completed using two NDMP servers. The data flows from one NDMP server to another NDMP server and then to the autochanger that is attached to the second NDMP server.

Caution – Three-party backups between an EMC NAS device and a NetApp server, or between a NetApp NAS device and an EMC server, must be recovered back to the original device.

FIGURE 11-3 demonstrates a three-party configuration.
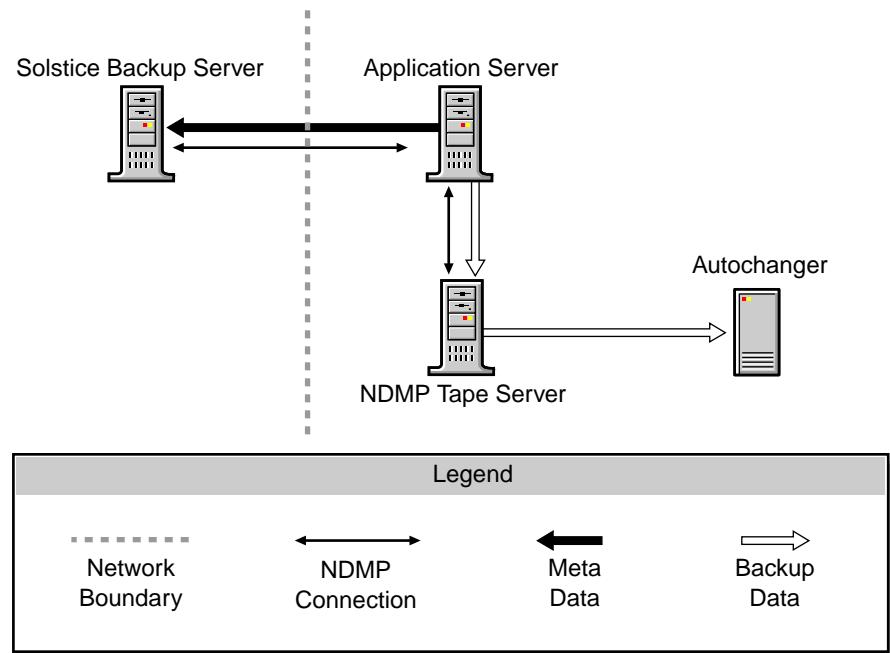
FIGURE 11-3 NDMP Three-Party Configuration

▼ To Perform a Three-Party Backup

To perform a three-party backup:

1. Configure the servers as shown in FIGURE 11-3 on page 451.

2. Install the Solstice Backup software on the server designated as your Solstice Backup server.

Note – You do not need to install any additional software on the NDMP servers.

---

3. Configure the Solstice Backup server:

   a. For the Name attribute of the Client Resource, specify the hostname of the NDMP server that contains the data to be backed up.

   b. For the Storage Node attribute of the Client Resource, verify that the destination server for the data being backed up is listed first.

The client backs up the data to each NDMP server or storage node in the order in which they appear in the attribute. If a successful backup cannot occur with the first NDMP server in the list, the client selects the next server listed in the Storage Node attribute.

For example, if the data you want to back up is located on a NetApp filer called venus and you want to back up the data to another NetApp filer called mars, you would list mars first and venus second in the Storage Node attribute. Once the data is successfully backed up to mars, the backup is finished. If there is a problem with mars, the data backs up to venus, since venus is the next server listed in the Storage Node attribute.

c. Complete the rest of the required Client resource attributes and configure the Solstice Backup server as you would for any other NDMP backup.

For instructions on configuring your Solstice Backup server, refer to "Setup and Configuration for NDMP Operations" on page 427.

Remote Backup

In a remote backup, the SnapImage Module software acts as an NDMP tape server and allows you to back up data from an NDMP disk server to a NetWorker server or storage node with the SnapImage Module software installed and tape drives attached. The SnapImage Module software is not actually performing the backup. The backup is controlled by the native NDMP utilities provided by the NAS box. This is not a block-level backup using the SnapImage Module software. The NetWorker server or storage node with the SnapImage Module software installed functions as a tape server, accepting the NDMP data stream and writing it to tape.

Note – Backing up data *from* the server or storage node with the SnapImage Module software *to* a tape drive attached to a NetApp Filer or EMC server is not currently supported.

FIGURE 11-4 on page 453 shows a possible configuration for backing up data on a NetApp, EMC, or Auspex server to a Solstice Backup server. In this configuration, the data flows from the NetApp, EMC, or Auspex server, over the LAN, and into the tape library directly attached to the Solstice Backup server. The SnapImage Module software is installed on the Solstice Backup server, acting as the NDMP tape server.
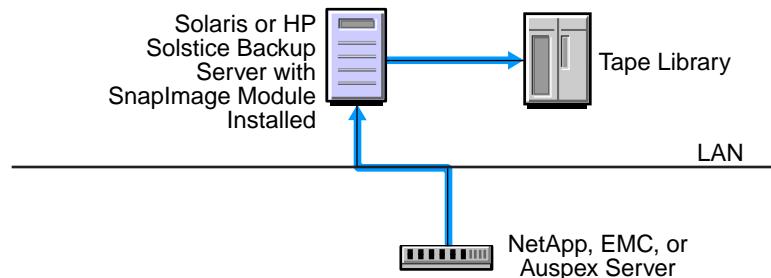
FIGURE 11-4 Performing a NetApp Backup to a Solstice Backup Server

FIGURE 11-5 shows a configuration where the data is backed up to a server or storage node with the SnapImage Module software installed. The data flows from the NetApp, EMC, or Auspex server, over the LAN, and into the tape library directly attached to the server or storage node. The backup is controlled by the native NDMP utilities provided by the NAS box. This is not a block-level backup using the SnapImage Module software. The server or storage node with the SnapImage Module software functions as a tape server, accepting the NDMP data stream and writing it to tape.
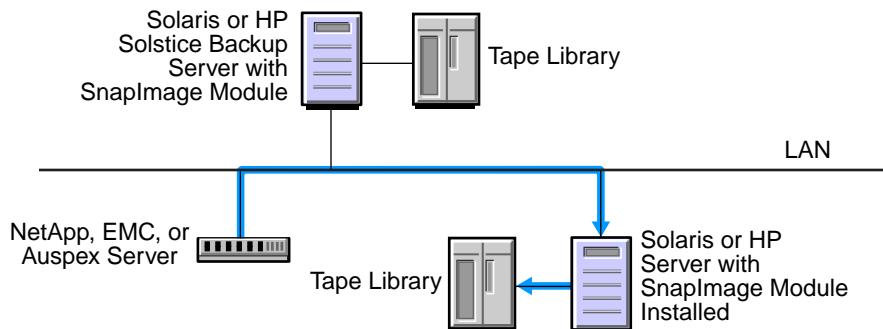
FIGURE 11-5 Performing a NetApp Backup to a SnapImage Module Server

▼ To Perform a Remote Backup

To perform a remote backup:

1. Configure the servers as shown in either FIGURE 11-4 on page 453 or FIGURE 11-5 on page 453.

2. Install the Solstice Backup software on the server designated as your Solstice Backup server.

3. Install the SnapImage Module software on the server that will act as a tape server:

   - If you are backing up the data to the Solstice Backup server, install the SnapImage Module software on that Solstice Backup server.

   - If you are backing up the data to a different NDMP server, install the SnapImage Module software on the server to which the data is being backed up.

   For instructions on installing the SnapImage Module software, refer to the *Legato NetWorker SnapImage Module Installation and Administrator's Guide*. Do not use the configuration instructions in the SnapImage Module guide. Instead, use the configuration instructions provided in this chapter.

4. To allow other servers to access the SnapImage server, run the ndmp_passwd command, located in /usr/ndmphome/bin.

   a. Use the following syntax:

      # ndmp_passwd -u *user*

      where *user* is the root or administrator name of the SnapImage server.

   b. When prompted, enter the password.

5. Configure the Solstice Backup server:

   a. For the Name attribute of the Client resource, specify the hostname of the NDMP server that contains the data to be backed up.

   b. For the Storage Node attribute of the Client resource, verify that the destination server for the data being backed up is listed first.

      The client backs up the data to each NDMP server or storage node in the order in which they appear in the attribute. If a successful backup cannot occur with the first NDMP server in the list, the client selects the next server listed for the Storage Node attribute.

      For example, if the data you want to back up is located on a NetApp filer called venus and you want to back up the data to a server with the SnapImage Module software installed called mars, you would list mars first and venus second for the Storage Node attribute. Once the data is successfully backed up to mars, the backup is finished. If there is a problem with mars, the data backs up to venus, since venus is the next server listed for the Storage Node attribute.

   c. Complete the rest of the required Client Resource attributes and configure the NetWorker server as you would for any other NDMP backup.

Performing NDMP Backups

A new command, `nsrndmp_save`, is used in NDMP operations. This command supports most of the same command-line arguments as `save(1m)`. An option within `nsrndmp_save`, `-T` *ndmp_type*, identifies the NDMP backup type.

The following sections describe how to perform an NDMP backup using the Solstice Backup Administrator program. You can also perform this task from the command line using the `nsrndmp_save` command. For information on this command, refer to the *Legato Command Reference Guide* or the `nsrndmp_save` man page.

Prerequisites

You can perform a test backup after you have:

- Enabled the NDMP Connection.
- Added the NDMP computer to a backup group.
- Selected a backup schedule with browse and retention policies.

Note – Incremental backups are not supported. If you select the incremental setting, the Solstice Backup server performs a full backup. However, you can schedule differential level backups to function like incremental backups. For example, a weekly backup schedule of full on day 1 and incrementals on days 2 through 7 is the same as a weekly backup schedule of full on day 1, level 1 on day 2, level 2 on day 3, level 3 on day 4, and so on.

- Selected a pool for your data.

Note – To prevent a non-NDMP volume from being recommended for an NDMP backup, create separate pools for NDMP and non-NDMP volumes.

- Installed an NDMP data module or enabled your application server for NDMP operations. Filesystems on computers without an NDMP data module cannot be backed up through the NetWorker NDMP Connection.
- Configured a non-NDMP device to back up the bootstrap information.

Next you will perform a test backup by starting a scheduled group backup.

▼ To Start a Scheduled Backup

Instructions for a stand-alone device are provided here. To use a device in an autochanger or silo, see Chapter 10.

Media Preparation

To prepare the media for the backup:

1. Insert a volume into the device you configured for Solstice Backup backups.

2. Select Media>Label or click the Label toolbar button to open the label window. The Label window displays the proposed label name, based on the label template you preconfigured for NDMP operations.

3. Select the Mount After Labeling attribute to automatically mount the volume following its labeling and click OK.

   The Solstice Backup software will label the volume and mount it in preparation for the test backup of your NDMP data.

NDMP Client Configuration

To configure your NDMP Client resource:

1. In the Solstice Backup Administrator program, select Clients>Client Setup to open the Clients window.

2. Select the NDMP client you created in the Clients attribute.

   The attribute values for your NDMP client are displayed.

3. Enter the save sets you want to back up in the Save Sets attribute and click Apply.

   To back up all the data on the NDMP computer, list all the filesystems you want to back up.

Note – Using savegroup probes or the All option for backups is not supported. You must manually list all the filesystems you want to back up.

Starting the Group Backup

To start the group backup:

1. In the Solstice Backup Administrator program, select Customize>Groups to open the Groups window.

2. Select the group that contains the NDMP client whose data you want to back up.

3. Select Start Now in the Autostart attribute and click Apply.

   The Solstice Backup server will now begin the group backup that includes your NDMP client.

While the backup is occurring, you can view the backup details.

Viewing the Group Backup Details

Follow these steps while the backup you just started is in progress.

1. In the Solstice Backup Administrator program, select Server>Group Control or click the Group Control toolbar button to open the Group Control window.

2. Select the Default entry in the Name attribute.

   This entry's status should indicate that it is in the process of being backed up.

3. Click Details to open the Group Control Details window.

   This window provides more detailed messages about the progress of the backup. When a backup is complete, lists of the successful, unsuccessful, and incomplete save sets are displayed.

Viewing Client File Index Information

To open the Indexes window after the backup has completed, select Clients>Indexes in the Solstice Backup Administrator program or click the Indexes toolbar button. Here you can view details of the NDMP client backup, as well as perform other tasks on your client file indexes. For general information regarding the management of the client file indexes, see "Online Indexes" on page 119.

If the backup did not complete successfully, refer to the troubleshooting information in Chapter A to determine the cause.

Performing NDMP Recoveries

A new utility, `nsrndmp_recover`, is used for NDMP operations. It supports most of the same command-line operations as `recover(1m)`, in addition to several others that specifically support NDMP.

The `nsrndmp_recover` command performs the following tasks:

- Connects to the appropriate Solstice Backup daemons
- Drives the recovery operation on the NDMP computer

If you retrieve data from the Solstice Backup Administrator program, the `nsrndmp_recover` command is selected automatically.

> **Caution –** The Save Set Recover options *Rename Recovered File* and *Discard Recovered File* are not supported by NDMP. *Current files are always overwritten*.

▼ To Perform a Destructive Save Set Recovery

A destructive save set recovery overwrites the entire contents of the filesystem. The following section describes how to perform a destructive save set recovery from the Solstice Backup Administrator program.

To perform a destructive save set recovery in the Solstice Backup Administrator program:

1. Select Save Set>Recover to open the Save Set Recover window.

2. Enter or select the NDMP client whose save set data you want to recover in the Client attribute.

   The save sets available for recovery are displayed in the Save Set attribute.

3. For the Save Set attribute, select the specific save set whose data you want to recover. You can select as many save sets as are listed in this attribute.

   When you select a save set, that save set's backup instances are then displayed in the Instances attribute.

4. Select the particular instance of a save set's backups in the Instances attribute.

5. Click Details if you want to view the details of the save set instance you selected, such as the backup date, number of backed-up files, save set ID number, and pool.

6. Click Recover to open the Save Set Recover Status window. Here you can customize the save set recovery in a variety of ways.

7. Specify the UNIX mount point in the Mount Point attribute using the following syntax:

```
filesystem_mount_point
```

For example, */UFS01/backup*.

8. Specify the raw device name of the destination in the Raw Device attribute using the following syntax:

```
destination_raw_device_name
```

For example, */dev/rdsk/c1t1d0s3*.

Note – If you do not specify the raw device name of the destination, the save set is restored and the files are overwritten, but the current filesystem structure is retained.

---

9. For the Paths to Recover attribute, use the default value.

10. Click Start.

You can also perform this task from the command line using the nsrndmp_recover command. Use the following syntax:

 # nsrndmp_recover -c *client* **-s** *server* **-r** *raw_device*
-S *ssid* **-m** *mount_point*

For example:

```
nsrndmp_recover -c venus -s mars -r /dev/rdsk/c1t1d0s3
-S 2485094913 -m /usr
```

For additional information on this command, refer to the *Legato Command Reference Guide* or the nsrndmp_recover man page.

For additional information on save set recoveries, including recovery options, see "Save Set Recovery" on page 306.

▼ To Perform a Directed Recovery

Solstice Backup software with NDMP now supports directed recoveries that allow you to recover backup data to a host other than the one used to back up the original data. The operating system on the target host (*hostname*) must be the same as the one used to back up the data.

For directed *Save Set* recoveries, do *one* of the following:

- In the Solstice Backup Administrator program:

    a. Select Save Set>Recover.

    b. Select the client and backup instance to be restored.

    c. Click Recover. The Save Set Recover Status window appears.

    d. Click Options. The Save Set Recover Options window appears.

    e. For the Relocate Recovered Data To attribute, enter the recovery information in the following format:

    *target_hostname***::**/*mount_point*

    For example, *molokai***::**/*usr*.

Note – For the mount point, specify the NAS box hostname. If you are using Data ONTAP 6.0, you might need to add a slash (/) after the mount point; for example, *target_hostname***::**/*mount_point*/.

---

    f. Click OK.

    g. Click Start in the Save Set Recover Status window.

- Using **nsrndmp_recover**, enter the following command:

    # nsrndmp_recover -c *client* **-m** *target_hostname***::**/*mount_point*
    -S *ssid*

    Note: For the mount point, specify the NAS box hostname. If you are using Data ONTAP ™ 6.0, you might need to add a slash (/) after the mount point; for example, *target_hostname***::**/*mount_point*/.

Note – In these command sequences, the host specified is the target host. If a host is not specified, the target is the original application server. The *mount_point* variable remains the previously defined mount point.

For directed *index* recoveries, do *one* of the following:

> **Note –** For the mount point, specify the NAS box hostname. If you are using Data ONTAP ™ 6.0, you might need to add a slash (/) after the mount point; for example, *target_hostname***::***/mount_point/*.

- Using the `nwrecover` program:

  a. Select Options>Relocate.

  b. For the Relocate attribute, enter:

     *target_hostname***::***/mount_point*

- Using the **recover** command, enter the following:

```
# relocate target_hostname::/mount_point
```

> **Note –** In these command sequences, the host specified is the target host. If a host is not specified, the target is the original application server.

Direct Access Restores

DAR is an optimized data recovery operation that allows you to retrieve data in the middle of a tape set without having to parse the tape set sequentially. This significantly reduces recovery time of large backups.

DAR is supported on the following platforms:

- NetApp Data ONTAP 6.1 and later
- Auspex 4Front NS2000 NetOS 3.0.1

You specify whether to use DAR by setting the application information variables DIRECT and HIST. For information about these settings, see "Specifying Application Information" on page 446.

DAR is supported in the following NDMP configurations:

- Local
- Three-party

A DAR backup is performed in the same manner as a regular NDMP backup. A DAR recovery is performed through the **recover** command or nwrecover.

Caution – To perform a recovery, use the **recover** command or nwrecover. You cannot use **nsrndmp_recover** for a recovery with DAR.

SNMP Module

The NetWorker Simple Network Management Protocol (SNMP) Module allows Solstice Backup servers to send notification messages to SNMP management agents.

SNMP is a protocol for network transactions that specifies the transfer of structured management information between SNMP managers and agents. This chapter describes the Solstice Backup server implementation of SNMP.

For detailed explanations of SNMP management operation, including how to configure your management software, refer to your network management software documentation.

The main sections in this chapter are:

- "SNMP Traps" on page 463
- "Configuring Legato NetWorker SNMP Notifications" on page 464
- "Configuring SNMP Management Software" on page 467

SNMP Traps

The Legato NetWorker SNMP Module uses *traps* to communicate Solstice Backup event notifications to SNMP management stations. A trap is an unsolicited notification sent from the SNMP agent (in this case, the Solstice Backup server) to the SNMP event manager.

The types of traps that the Solstice Backup server sends are determined when the Legato NetWorker SNMP notification is configured within your Solstice Backup server. Typical traps include warnings, critical errors, and other messages from the Solstice Backup server. For instructions on configuring the Legato NetWorker SNMP notification, see "Configuring Legato NetWorker SNMP Notifications" on page 464.

Additionally, your SNMP-enabled network management software must be configured to accept traps from the Solstice Backup server. Depending on the configuration of your network management software, traps received from the Solstice Backup server may be displayed in the management software's graphical interface, or may be logged into a file or processed in another manner.

Configuring Legato NetWorker SNMP Notifications

Solstice Backup software provides notifications to a variety of resources about Solstice Backup server events. The Legato NetWorker SNMP Module is one of those resources. The module then forwards the notification to your SNMP management software using the `nsrtrap` program. When you configure the SNMP notification, you include the IP address or hostname of your SNMP management server, along with other optional `nsrtrap` command line switches, such as the SNMP community and the trap type.

In order to configure the Legato NetWorker SNMP notification, you must first enable the Legato NetWorker SNMP module. For information about enabling the Legato NetWorker SNMP module, see the Solstice Backup Roadmap.

Command-Line Options for **nsrtrap**

The Legato NetWorker SNMP Module uses the `nsrtrap` program to communicate SNMP traps from the Solstice Backup server to your SNMP management software. TABLE 12-1 lists the command-line options that can be included in the Action attribute when the SNMP notification is configured:

TABLE 12-1   Command-Line Options for nsrtrap *(1 of 2)*

Option	Description
–c *community*	This option allows you to specify the SNMP community that is authorized to receive traps from the Solstice Backup server. SNMP communities are configured on your SNMP server. The default setting for this option is Public, meaning that the public community can receive traps from the Solstice Backup server. For security purposes, system administrators often customize their SNMP servers to limit the communities from which their server will accept traps. If a community other than Public has been configured on your SNMP server, you should include the appropriate community name using this option when you configure your SNMP notification.

TABLE 12-1 Command-Line Options for nsrtrap *(2 of 2)*

Option	Description
-t *trap_type*	This option is used to set the type of trap the Legato NetWorker SNMP Module sends to your SNMP server. The default setting is 6, which means that this is an "enterprise-specific" trap. Because traps sent by the Solstice Backup server are notifications (error messages, for example), the default setting is normally correct and should not be changed. This option should be used only if you intend to send a specific trap other than a normal Solstice Backup notification.
-s *specific_type*	This option is a generic setting that can be used to identify the type of trap being sent by the Solstice Backup server. This option can be set to any integer value and may be used in conjunction with different SNMP notifications to distinguish different traps coming from the Solstice Backup server. For example, you can create multiple SNMP notifications: one for critical messages, another for warnings, and another for other events or priorities. You can then use the -s option to differentiate the various notifications so that your SNMP management software can determine which type of trap is being sent. You could create one notification called Critical SNMP Notification, and include the -s option in the Action attribute: **/usr/sbin/nsr/nsrtrap -s 1** *host* With this setting, your SNMP management software can be configured to recognize that Solstice Backup traps with the specific trap type of **1** are critical messages. Additional SNMP notifications can have other settings for the -s option to further differentiate various traps from the Solstice Backup server.
-v	This option is used to set the output mode to verbose. In verbose mode, nsrtrap echoes the community, trap type, specific trap type, and the hostname or IP address to the command line.

▼ To Modify the Preconfigured Legato NetWorker SNMP Notification

The Solstice Backup server has a preconfigured SNMP notification that can be modified if necessary. The only modification that can be made to this notification is to add or remove command-line options to the Action attribute.

To modify the preconfigured notification request:

1. In the Solstice Backup Administrator program, select Customize>Notifications to open the Notifications window.

2. From the View menu, select Details to display the event and priorities for this notification.

3. Highlight the SNMP Notification.

4. To add command-line options for nsrtrap, such as the SNMP community, enter any necessary options to the command listed in the Action attribute. For more information about command-line options, see "Command-Line Options for nsrtrap" on page 464.

   The events and priorities associated with the preconfigured SNMP notification cannot be modified. If you want to set different events and priorities for your SNMP notification, see "To Create Legato NetWorker SNMP Notifications" on page 466.

▼ To Create Legato NetWorker SNMP Notifications

To create additional Legato NetWorker SNMP notifications:

1. In the Solstice Backup Administrator program, select Customize>Notifications to open the Notifications window.

2. From the View menu, select Details to display the event and priorities for this notification.

3. Click Create.

4. For the Name attribute, enter "SNMP notification request."

5. For the Action attribute, enter the path to your Solstice Backup executables, along with nsrtrap *network_management_station*, where *network_management_station* is the DNS name or IP address of the host on which your SNMP management software is running.

   Command-line options for nsrtrap, such as the SNMP community, should be included in this attribute if necessary. For more information about command-line options, see "Command-Line Options for nsrtrap" on page 464.

   For example, the Action attribute might read:

```
/usr/sbin/nsr/nsrtrap SNMPhost -c sbu
```

6. Select the events and priorities to be communicated by the notification.

Note – The events and priorities cannot be modified after the notification has been created.

---

7. Click Apply.

Configuring SNMP Management Software

In order for your SNMP management software to accept traps sent by Solstice Backup servers, it must be configured to recognize the traps. Instructions for configuring SNMP management software to accept specific traps vary depending on the type of management software you are using.

For specific instructions on configuring the types of traps that the software will accept, refer to your SNMP management software documentation.

Solstice Backup SMI Network Management Private Enterprise Code

When configuring your management software to accept traps, you must also indicate the specific type of trap that you want it to accept. This is done using the SMI Network Management Private Enterprise Code that applies to the specific network application that will be sending traps to your management software. The Private Enterprise Code for the Solstice Backup server is 160 (the complete code is actually .1.3.6.1.4.1.160).

Receiving Traps in Your SNMP Network Management Software

Typically, once your network management software is configured to accept traps from Solstice Backup servers, icons representing each Solstice Backup server are displayed on your network management console. You can then configure your

management software to determine what it should do when it receives a trap from a Solstice Backup server. Using your network management software, you may want to:

- Configure the way in which the management software indicates that a trap was received (for example, the Solstice Backup server icon may blink or change color).

- Track pending, alert, and other configured messages.

- Separate traps into event categories, such as Error Events, Status Events, Threshold Events, Configuration Events, Application Alert Events, or All Events. For information on setting up SNMP trap templates, refer to your network management software documentation.

You may also want to create additional SNMP notification schemes in the Solstice Backup Administrator program that have different priorities and events. You can use the -s *specific-type* command-line option for nsrtrap so that your SNMP management software can differentiate the traps sent by the various notification schemes. For more information about setting the -s *specific-type* option, see "Command-Line Options for nsrtrap" on page 464.

Notes and Troubleshooting Tips

If you have a problem with Solstice Backup software, or if the product does not work the way you expect, use the information in this appendix to diagnose your problem.

This appendix includes the following sections:

Information to Gather Before You Call Technical Support

If the solutions in this appendix do not solve the problem, be prepared to provide the following information when you call Sun Technical Support:

- The software release of the Solstice Backup component.
- The version of operating system that you are running. You can determine this with the **uname -a** command.
- Your hardware configuration.
- Information on your devices and other SCSI IDs. Use the **/etc/LGTOuscsi/inquire** command as root to obtain the required information.

- If you are using an autochanger, the type of connection (SCSI or RS-232). Also, provide the version of the autochanger driver you are using. You can determine this from the output of **pkginfo -x SUNWsbus2**.
- You should also be able to relate the following:
  - How to reproduce the problem
  - The exact error messages
  - How many times you have seen the problem
  - Whether the Solstice Backup command was successful before you made any changes and, if so, the changes you made

Solstice Backup Backup and Recover

This section explains how to troubleshoot various problems you might encounter with Solstice Backup backup and recover operations.

Multiple `nsrexec` Processes

You might notice multiple instances of the `nsrexec` process running on the Solstice Backup server during a backup. This is normal behavior and happens because one `nsrexec` process is spawned for each active save stream.

Checking the Solstice Backup Daemons

If you have trouble starting Solstice Backup programs, the daemons might not be running properly. To determine whether the required daemons are running, enter the following command:

```
# ps -ef | grep nsr
```

or

```
# ps -ax | grep nsr
```

You should receive a response similar to the following:

```
12217 ?          S  0:09 /usr/sbin/nsr/nsrexecd -s jupiter

12221 ?          S  2:23 /usr/sbin/nsr/nsrd

12230 ?          S  0:00 /usr/sbin/nsr/nsrmmdbd

12231 ?          S  0:01 /usr/sbin/nsr/nsrindexd

12232 ?          S  0:00 /usr/sbin/nsr/nsrmmd -n 1

12234 ?          S  0:00 /usr/sbin/nsr/nsrmmd -n 2

12235 ?          S  0:00 /usr/sbin/nsr/nsrmmd -n 3

12236 ?          S  0:00 /usr/sbin/nsr/nsrmmd -n 4

12410 pts/8     S  0:00 grep nsr
```

If the response indicates that the daemons are not present, start the Solstice Backup daemons with the following command:

```
# /etc/init.d/networker start
```

Unsynchronized Client Clock Errors

This situation affects computers with Solstice Backup client software, release 5.0 or later. If the setting for the system clock on a Solstice Backup client differs from that of the Solstice Backup server by more than one minute, the following message appears during a scheduled backup that invokes the savegrp program:

```
Warning: unsynchronized client clock detected
```

If the gap in synchronization is less than two minutes, these messages appear, but the schedule is still honored. To avoid receiving these messages, keep the server and client clocks synchronized to within one minute of each other.

Backups Fail to Stop

Attempting to stop the backup process by clicking Stop in the Group Control window should stop the process for all clients in the selected group. However, sometimes a client is missed and messages appear indicating the server is still busy.

To resolve the problem:

1. Determine the group that is currently being backed up by looking at the messages displayed in the Solstice Backup Administrator program.

2. Determine the process ID for the currently running save process using *one* of the following commands:

```
# ps -ef | grep save
```

```
# ps -ax | grep save
```

If the group status shows that the save processes are running, but the associated savegrp process is not running, perform one of these options:

- Stop the conflicting group from running by clicking Stop in the Group Control window. For more information, see

- Shut down and restart the Solstice Backup daemons. To do this:

   1. Shut down the Solstice Backup daemons using the nsr_shutdown -a command.

   2. Restart the Solstice Backup daemons using the nsrd command.

- Stop the save process for each process ID, using the kill -9 pid command.

Save Set Cutoffs for 5.x Clients

The Save Set Cutoff option is no longer supported in Solstice Backup 6.0 or later releases. When Solstice Backup clients earlier than 6.0 back up to a Solstice Backup 6.0 or later server, the server will respect the save set cutoff value. If both the server and the client are Solstice Backup 6.0 or later releases, then the save set cutoff feature is not present.

Licensing-Based Errors with Release Clients Earlier Than 5.0

Clients running Solstice Backup software earlier than release 5.0 do not have "self-ID" capabilities and are unable to report their OSs to the server at backup time. These clients are all listed under pre-5.0 client types when `nsrlic -v` is run. Since the server does not know what OS the client is running, it will always use a workstation license, even if the client is on a Windows NT or UNIX server. However, such non-ID clients have no effect on which ClientPak licenses are required to allow Solstice Backup clients to function; the appropriate licenses are still required.

In the following cases, Solstice Backup backups are rejected because of licensing enforcement, and an error message is generated:

- Backups are rejected if the client type is not allowed by the enabler code, or if a ClientPak enabler is not present. For example, the following error message might be sent for a Solaris server with an unlicensed Windows NT client:

```
RAP error, ClientPak for NT does not exist in the system.
```

- Backups are rejected if no more licenses are available for that type of client. For example, you might receive an error message similar to this:

```
RAP error, Too many clients. Maximum is 2
```

Solstice Backup clients released before Solstice Backup 5.0 identify themselves to the Solstice Backup server at backup time.

Check Failure of Client File Indexes

Each time the Solstice Backup server starts, it uses `nsrck -ML1` to perform a level 1 consistency check on the client file indexes. In some circumstances, this consistency check will not detect corruption in the client file indexes. If you believe an index may be corrupt, you should run a higher level check on the index, for example:

```
# nsrck -L5
```

If the index is still corrupt, recover the index using the procedure outlined in "Recovering Client File Indexes" on page 316.

No Notification of Client File Index Size Growth

The Solstice Backup server does not notify you when a client file index is getting too large. Monitor the system regularly to check the size of client file indexes. For information on how to manage the Solstice Backup client file indexes, see "Online Indexes" on page 119. For more information, refer to nsrls, nsrck, and nsrim in the *Legato Command Reference Guide* or the nsrls, nsrck, and nsrim man pages.

Media Position Errors Encountered When Auto Media Verify Is Enabled

When you enable Auto Media Verify for a pool, the Solstice Backup server verifies the data written to volumes from the pool during the save. This is done by reading a record of data written to the media and comparing it to the original record. Media is verified after the server finishes writing to the volume, which might occur when a volume becomes full or when the server no longer needs the volume for saving data.

To verify media, nsrmmd must reposition the volume to read previously written data. It does not always succeed in the first attempt. These warning messages appear in the message display in the Solstice Backup Administrator program:

```
media warning: /dev/rmt2.1 moving: fsr 15: I/O error

media emergency: could not position jupiter.007 to file 44, record
16
```

No action is required. The Solstice Backup server continues to attempt to find the proper position. If the server can find the correct position, media verification succeeds and a successful completion message appears.

```
media info: verification of volume "jupiter.007" volid 30052
succeeded.
```

In this case, ignore the earlier messages because they only indicate that the server had problems finding the desired position on the media.

If media verification fails, attempt the following:

■ Reset the device.
■ Verify the configuration of the device.

- Verify that the media can be recognized.
- Verify that the device is functioning properly.

PACKET RECEIVE BUFFER and NO ECB Counters Increase

When your server is waiting for a tape to be mounted or is in the process of changing an autochanger volume, the PACKET RECEIVE BUFFER and NO ECB counters increase on a NetWare client.

To resolve this problem, use the `nsr_shutdown` command to shut down the Solstice Backup server.

Next, restart the Solstice Backup server manually using the following command:

```
# /etc/init.d/networker start
```

The `scanner` Program Marks a Volume Read-Only

When you use the `scanner` program to rebuild the index of a backup volume, the `scanner` program marks the volume as read-only.

This is a safety feature that prevents the last save set on the backup volume from being overwritten. To write to the media without marking it read-only, use the `nsrmm -o` command:

```
# nsrmm -o notreadonly volume_name
```

Index Recovery to a Different Location Fails

If you attempt to recover indexes to a directory other than the one where they were originally located, you will receive the following error message:

```
WARNING: The on-line index for `client_name' was NOT fully recovered.
There may have been a media error. You can retry the recover, or
attempt to recover another version of the index.
```

Do not attempt to recover the indexes to a different directory. After the indexes have been recovered to their original location, you can move them to another directory. To move the indexes, invoke the following command as root from within the /nsr/index directory:

```
# /usr/lib/nsr/uasm -s -i client_index_directory_name | (cd target_directory;
/usr/lib/nsr/uasm -r)
```

Potential Cause for Client Alias Problems

If you encounter any of the following situations, a client alias problem might be the cause:

- You receive the following error message:

```
No client resource for client_name
```

- A client machine always performs full backups, regardless of the level of the scheduled backup.
- It appears that automatic index management, as set up in the browse and retention policies, does not occur. This is indicated by the filesystem containing the indexes continuously increasing in size.
- In /nsr/index (the directory that contains the indexes) there are two directories for the same client using two different client names.

A client alias change is needed for the following situations:

- Machines that have two or more network interfaces
- Sites that mix short and fully qualified hostnames for the same machines; for example, mars and mars.jupiter.com
- Sites using both YP (NIS) and DNS

Edit the Client resource for clients with this problem. Add all network names for this host to the Aliases attribute.

Caution – Do not put aliases that are shared by other hosts on this command line.

Illegal Characters to Avoid in Configurations

When naming label templates, directives, groups, policies, and schedules, the following characters are not allowed:

```
/ \ * [ ] ( ) $ ! ^ ' ; ` ~ < > & | { _
```

In addition, the following characters cannot be used in label templates because they are reserved for use as field separators:

- colon (:)
- hyphen (-)
- period (.)
- underscore (_)

The `scanner` Program Requests an Entry for Record Size

If you use the scanner program with the -s option but without an -i or -m option, and you receive the message:

```
please enter record size for this volume ('q' to quit) [xx]
```

the number in the bracket [xx] is the entry from the last query.

The scanner command always rewinds the tape and reads the volume label to determine the block size. If the volume label is corrupted or unreadable, you see a message prompting you to enter the block size (in kilobytes).

Type in the block size; it must be an integer equal to or greater than 32. If you enter an integer that is less than 32, you receive the following message:

```
illegal record size (must be an integer >=32)
```

Failed Recover Operation Directly after New Installation

If you attempt to start the `nwrecover` program immediately after installing Solstice Backup software for the first time on your system, you receive the error message:

```
nwrecover: Program not found.
```

To save disk space, Solstice Backup software delays the creation of the client index until the first backup is completed. The `nwrecover` program cannot recover data until the client index has entries for browsing. To avoid the problem, perform a Solstice Backup backup on the client.

Client File Index Is Missing Message

If you attempt to recover the client file indexes with the `scanner -i` without first using `nsrck -c` to create a new index, you might encounter a message similar to the following:

```
scanner: File index error, file index is missing.

Please contact your system administrator to recover or recreate
the index.
(severity 5, number 8)

scanner: write failed, Broken pipe
scanner: ssid 25312: scan complete
scanner: ssid 25312: 91 KB, 13 file(s)
scanner: done with file disk default.001
```

The `scanner` program must have a file index to rebuild from before it can proceed. To resolve the problem, always run the `nsrck` program with a `-c` option to create an index before you run `scanner -i` to rebuild the client file indexes from the backup media.

Recovering Files from an Interrupted Backup

If you terminate a backup by killing the Solstice Backup daemons, you cannot recover the files because the media database is not updated when the daemons die. Consequently, the Solstice Backup server does not know on which volumes the requested files reside.

Backup of a New Client Defaults to a Level Full

The first time you back up a new client, you receive the following message:

```
client: save point: There are no save sets in the media database;
performing a full backup
```

This message indicates that the specified save set has not been previously backed up. If the save set had been backed up, the save set would appear in the media database.

Before an incremental or level backup can be performed on a save set, it must first be backed up in its entirety.

If the save set was previously backed up, this message might appear because:

■ The clocks on the client and server are not synchronized.
■ The savegrp session begins before midnight and ends after midnight.

Renamed Clients Cannot Recover Old Backups

The Solstice Backup server maintains a client file index for every client it backs up. If you change the name of the client, the index for that client is not associated with the client's new name and you cannot recover files backed up under the old client name.

To recover data that was backed up using the old client name, you should perform a directed recovery by directing data saved under the old client name to the new client. For information about performing directed recoveries, see "Directed Recovery" on page 299.

The `savegrp` Command-Line Program Cannot Be Stopped from GUI

If you start the `savegrp` command at the command prompt and attempt to stop the backup from the Solstice Backup Administrator Program, the following message appears:

```
Only automatically started groups that are currently running can
be stopped
```

To stop the command, you must first determine the process ID of the `savegrp` process, and then use the UNIX `kill` command to kill the process.

▼ To Stop the `savegrp` Command Line Program

To stop the `savegrp` command-line program:

1. Determine the `savegrp` process ID using the `ps` command:

```
# ps -ef | grep savegrp
```

The process ID is displayed in the following format:

```
root  process_id 4769   0 13:53:10 pts/2     0:00 savegrp myGroup
```

2. Stop the `savegrp` process using the `kill` command:

```
# kill process_id
```

Disk Label Errors

If you receive the error message:

```
No disk label
```

you might have incorrectly configured a nonoptical device as an optical device within the Solstice Backup server. Verify that the Media Type attribute in the Devices resource matches the expected media for your device, and make corrections if needed.

Cannot Print Bootstrap Information

If your server bootstraps are not printed, you might need to enter your printer's name as a hidden attribute in the Groups resource.

1. Select View>Details to display hidden attributes.

2. Enter the name of the printer where you want the bootstraps to be printed in the Printer attribute of the Groups resource.

Server Index Not Forced

If your Solstice Backup server belongs to a group that is not enabled or does not belong to any group, the savegrp program does not back it up, because it is not necessary to force a server index backup.

The information to recover server indexes is stored in the media database.

Caution – If your datazone's planning has depended on savegrp forcing a server index backup that has never been explicitly backed up, note that savegrp no longer performs a server index backup for servers not in an active group.

Copy Violation

If you installed Solstice Backup software on more than one server and used the same Solstice Backup enabler code for them all, you receive messages similar to the following in your save group completion mail:

```
--- Unsuccessful Save Sets ---

* mars:/var save: error, copy violation - servers 'jupiter' and
'pluto' have the same software enabler code, 'a1b2c3d4f5g6h7j8'
(13)

* mars:/var save: cannot start a backup for /var with NSR server
'jupiter'

* mars:index save: cannot start a backup for /usr/nsr/index/mars
with NSR server 'jupiter'

* mars:index save: cannot start a backup for bootstrap with NSR
server 'jupiter'

* mars:index save: bootstrap save of server's index and volume
databases failed
```

To successfully rerun the backup, you must issue the nsr_shutdown command on each server, remove the Solstice Backup software from the extra servers, and then restart the Solstice Backup daemons on the server where you want the backups to go.

Xview Errors

If you receive the following error message when you attempt to start the graphical administration interface with the Solstice Backup Administrator program from a client machine, it means that the client is not authorized to display Solstice Backup GUIs:

```
Xlib: connection to "mars:0.0" refused by server

Xlib: Client is not authorized to connect to Server

Xview error: Cannot open display on window server: mars:0.0 (Server
package)
```

To correct the situation, configure the client to display Solstice Backup GUIs:

1. From the client machine, invoke the xhost command:

```
# xhost server_name
```

2. Remotely log in to the Solstice Backup server and issue the setenv command at the prompt.

```
# setenv DISPLAY client_name:0.0
```

For command shells other than csh enter:

```
# DISPLAY=client_name:0.0

# export DISPLAY
```

Converting Sparse Files to Fully Allocated Files

The Solstice Backup server determines that files are sparse or holey by comparing the allocated blocks with the byte size. Use the ls -ls command to compare blocks. If the allocated blocks do not account for the size of the file, the file is considered to be sparse and is saved using an algorithm that replaces long strings of zeroes with "holes" in the recovered file.

Some files that were not sparse when saved might be recovered as sparse. Oracle databases are particularly susceptible to this problem because they are zero-filled, fully allocated files and are not sparse.

At this time, the workaround for this situation is to use the `cp` command to copy the file after recovery:

    # **cp** *recovered_filename zero_filled_filename*

This converts a sparse file to a fully allocated file.

Note – If you choose this workaround, you must have enough free disk space to accommodate a duplicate of each sparse file you choose to copy.

Client/Server Communications

Many of the problems that Solstice Backup users report when they set up and configure Solstice Backup software are problems with the communications in their networks. This section contains a procedure for testing the communications in a network.

The information provided here is for UNIX platforms only. If your Solstice Backup setup includes other platforms, refer to *Legato Technical Bulletin 299*, which is included in the `bulletins.pdf` file that shipped with your Solstice Backup binaries.

▼ To Troubleshoot IP Errors

If any of the following error conditions occurs, you probably have an IP naming problem:

- RPC errors
- Unknown host messages
- Failure contacting the portmapper
- Connection failures or timeouts
- Program unexpectedly exiting
- Refused connections
- RCMD to active client failures
- Name-to-address translation failures

- Program not registered messages
- Solstice Backup services or daemons failing to start
- Solstice Backup services or daemons failing to remain active
- Invalid path messages

To troubleshoot IP errors:

1. Document the steps you take and the results, especially error messages, in case you need to contact Sun Technical Support. This enables you to e-mail or fax the exact steps and error message text directly to Sun.

2. Set up a hosts table for each Solstice Backup client and Solstice Backup server. For more information, see "To Set Up a Hosts Table" on page 485.

3. Disable other name servers to simplify testing. For more information, see "Troubleshooting Hostname Resolution Problems" on page 486.

4. Use ping to establish basic connectivity. For more information, see "To Use ping to Verify Network Connections" on page 487.

5. Use rpcinfo to verify that sessions can be established and that portmapping is correct. See "To Use rpcinfo to Verify That Sessions Can Be Established" on page 488 for more information.

▼ To Set Up a Hosts Table

It is recommended that you troubleshoot IP problems using only the hosts table. Troubleshooting using only the hosts table does not mean you cannot use your name service, for example, DNS, with Solstice Backup software. Run tests using only the hosts table to determine whether you have the software installed correctly. After you know the software works with the hosts table, you can enable the named server you are using.

To set up a hosts table:

1. On the Solstice Backup client, list the client and the Solstice Backup servers to which it connects, for example:

```
127.0.0.1 localhost loopback

123.456.789.111 client client.domain.com

123.456.789.222 server server.domain.com
```

2. On the Solstice Backup server, list the Solstice Backup server and all of its clients, for example:

```
127.0.0.1 localhost loopback
```

123.456.789.111 server server.domain.com

123.456.789.222 client client.domain.com

3. Use the guidelines in "To Use ping to Verify Network Connections" on page 487 to ensure the highest success rate for parsing a hosts table within any operating system.

Recommendations for hosts table configuration:

- Do not use blank lines in the body of your hosts table.
- The end of the hosts table should always contain a blank line.
- The first unremarked entry should always be the loopback line in the exact order and format shown in Step 1 and Step 2.
- The last character of each unremarked line should be a space, not a carriage return.

On UNIX platforms, the hosts table resides in the /etc/hosts directory.

Troubleshooting Hostname Resolution Problems

To simplify the troubleshooting of host name resolution problems, try disabling services like DNS, WINS, and DHCP. If you have hostname resolution problems, first configure the hosts table for each computer, and then test your backups.

Some common problems you can encounter with DNS, WINS, and DHCP services include:

- The DNS is not configured with a reverse lookup table.
- The clients are configured with the wrong IP addresses for DNS or WINS servers.
- The DHCP services do not properly update the WINS server with new addresses.

You do not need to disable DNS for your entire network, just for the initial setup of the Solstice Backup clients and the Solstice Backup server you want to test. Only disable the ability of a client to obtain IP naming information from a DNS server. Typically, you do not need to disable the DNS server itself.

To disable the DNS server on most UNIX platforms, rename the file */etc/resolv.conf* and reboot.

Instead of renaming `resolv.conf`, you can set up the IP name search order so that the hosts table is searched before DNS.

▼ To Set the IP Name Search Order

To set the IP name search order:

1. Edit the */etc/nsswitch.conf* file and verify that the */etc/resolv.conf* file exists.

2. Set the host file to be first in the search order, with DNS second and NIS last, for example:

```
hosts: files [NOTFOUND=continue] DNS [NOTFOUND=continue] nis
```

You can also set the NSORDER environment variable. For directions specific to your version of AIX, refer to Info Explorer.

▼ To Use `ping` to Verify Network Connections

After you create the hosts table, test it using `ping`.

On the Solstice Backup client, `ping` the following:

- Client short name (hostname) from the client
- Client long name (hostname plus domain information) from the client
- Client IP address from the client
- Server short name from the client
- Server long name from the client
- Server IP address from the client

The following example shows how to `ping` the client short name and client long name from a Solstice Backup client called *mars* in the *jupiter* domain:

```
# ping mars

# ping mars.jupiter.com
```

On the Solstice Backup server, `ping` the following:

- Server short name from the server
- Server long name from the server
- Server IP address from the server

- Client short name from the server
- Client long name from the server
- Client IP address from the server

▼ To Use `rpcinfo` to Verify That Sessions Can Be Established

If `ping` is successful and backup problems still exist, you can also test with `rpcinfo`. Because Solstice Backup software relies heavily on mapping of ports, use `rpcinfo` to test the operation of the portmapper. Using `ping` tests the connection up to the network layer in the Open System Interconnection (OSI) model; `rpcinfo` checks for communication up to the session layer.

Use the same tests with `rpcinfo` as with `ping`.

For `rpcinfo` to be used successfully, the machine whose hostname you enter on the command line must have a portmapper running. In most cases, Sun portmappers are compatible with fully functional portmappers from other vendors (this is called a third-party portmapper). If you are using a product that provides its own portmapper, we recommend not loading the third-party portmapper until you have verified that the Solstice Backup software works with the rest of your environment. This process lets you test portmapper compatibility without adding other unknowns.

The `rpcbind` daemon must be running. The **rpcinfo** utility is part of the operating system.

The syntax for using `rpcinfo` to display ports using TCP is:

```
# rpcinfo -p hostname
```

Substitute the long name and short name for the variable *hostname*, just like for `ping`.

View other `rpcinfo` command-line options by typing `rpcinfo` at the command line. Notes on the `rpcinfo` command and its error messages are available in the UNIX man page for `rpcinfo`. Repeat `rpcinfo` using all the locations and all the iterations listed in this document for `ping`.

When `rpcinfo` runs successfully, the output is a list of port numbers and names as shown in the following example:

```
rpcinfo for mars

program vers proto    port

100000    2   tcp     111  portmapper

100000    2   udp     111  portmapper

390103    2   tcp     760

390109    2   tcp     760

390110    1   tcp     760

390103    2   udp     764

390109    2   udp     764

390110    1   udp     764

390113    1   tcp    7937

390105    5   tcp     821

390107    4   tcp     819

390107    5   tcp     819

390104  105   tcp     822
```

Replacement for Functional Use of `portmapper` and `rpcbind`

The `nsrexecd` process now combines the functionality of `portmapper` and `rpcbind`. The parent `nsrexecd` spawns a child `nsrexecd` that handles the `portmapper` functions. As a result, you will see more than one instance of the

nsrexecd process. In the following example, the instance of nsrexecd running under the process ID (PID) 20415 is the parent and the instance of nsrexecd running under the PID 20416 is the child:

```
jupiter# ps -ef | grep nsr

root 20429 204200 11:23:59 ? 0:00 /usr/sbin/nsr/nsrmmd -n 1
root 20415 1 0 11:23:49 ? 0:01 /usr/sbin/nsr/nsrexecd -s jupiter
root 20416 204150 11:23:49 ? 0:01 /usr/sbin/nsr/nsrexecd -s
jupiter
root 20428 204200 11:23:57 ? 0:00 /usr/sbin/nsr/nsrindexd
root 20420 1 0 11:23:53 ? 0:02 /usr/sbin/nsr/nsrd
root 20446 203490 11:27:22 pts/6 0:00 grep nsr
root 20427 204200 11:23:55 ? 0:00 /usr/sbin/nsr/nsrmmdbd
```

This change allows the Solstice Backup processes to function in a more secure manner that does not utilize ports or protocols that are commonly used by other programs.

▼ To Verify Firmware for Switches and Routers

If you are using switches or routers from any vendor, make sure that the switch or router firmware is dated after August 1995 (wherever they exist on your network) to ensure that RPC (Remote Procedure Call) traffic is handled properly. Most of the switch and router vendors have significantly improved their handling of RPC traffic since August 1995.

Naming Requirements

Solstice Backup UNIX clients, release 4.2 and later, use the servers file in the /nsr/res subdirectory to determine whether a Solstice Backup server is authorized to back up the client's data. If you do not have the servers file, you can create it in /nsr/res using your preferred editor.

Make sure the `servers` file on a client contains both the short name and the long name of the server you want to use to back up that client's data. For example, the `servers` file on a Solstice Backup client would contain the following names for a Solstice Backup server named `mars` in the `jupiter.com` domain:

```
mars

mars.jupiter.com
```

In the Alias attribute of the Clients resource, list both the short name and the long name, plus any other applicable aliases for each client.

Binding to Server Errors

Solstice Backup architecture follows the client/server model, where servers provide services to the client through the RPC. These services reside inside of long-lived processes known as daemons.

For clients to find these daemons, register the daemons with a registration service. When the daemons start, they register themselves with the registration service provided by the portmapper.

Solstice Backup servers provide a backup and recovery service. The servers receive data from clients, store the data on backup media, and retrieve it on demand. If the Solstice Backup daemons are not running and a Solstice Backup service is requested, you receive the following messages in your save group completion mail:

```
Server not available

RPC error, no remote program registered
```

These messages indicate that the Solstice Backup daemons `nsrd`, `nsrexecd`, `nsrindexd`, `nsrmmd`, and `nsrmmdbd` might not be running.

To restart the daemons, log in as root and enter the following command at the prompt:

```
# /etc/init.d/networker start
```

Saving Remote Filesystems

You might receive the following error messages in your save group completion mail when a backup for a remote client fails:

```
Host hostname cannot request command execution

hostname: Permission denied
```

The first message means that the `nsrexecd` daemon on the client is not configured to allow the server to back up its files. The second message means that the `nsrexecd` daemon is not currently running on the client.

To resolve these problems, make sure that the `nsrexecd` daemon is running on the client, and that the server's hostname is listed in the boot-time file. The boot-time file is automatically generated before the installation script is completed, and takes your responses to the query for the names of all the servers, in order of precedence, that can contact a client for backups. TABLE A-1 lists the location for the boot-time file. For detailed information about the nsrexecd daemon, refer to the `nsrexecd(1m)` man page.

TABLE A-1 Boot-Time File Locations

Operating System	Boot-Time File
Solaris	`/etc/init.d/networker`
SunOS 4.1.x	`/etc/rc.local`

Remote Recover Access Rights

You can control client recover access by configuring the Client resource. The Remote Access list displays the user names that have recover access to the client's save sets. You can add or remove user names depending on the level of security the files require.

The following users have permission to recover any files on any client, regardless of the contents of the Remote Access list:

- Root
- Operator
- Member of the operator group

Other users can only recover files for which they have read permission, relative to the file mode and ownership at the time that the file was backed up. Files recovered by a user other than root, operator, or the operator group are owned by that user.

Networker Archive and Retrieve

This section explains how to troubleshoot various problems you might encounter during archiving and retrieval with the Networker Archive Module.

Remote Archive Request from Server Fails

If you cannot perform a remote archive request of a computer from the Solstice Backup server, the archive client's user name (for example, `root`) might not be listed in that client's Archive Users attribute in the Clients resource.

You can also grant Solstice Backup administrator privileges for `root@client_system` in the Administrator attribute in the Server resource. Granting administrator privileges creates a potential security issue, because Solstice Backup administrators can recover and retrieve data owned by other users on other clients.

Multiple Save Sets Appear as a Single Archive Save Set

When you combine multiple save sets in an archive, such as `/home` and `/usr`, they end up in a single archive save set, which appears as a slash (/) in the Archives list in the `nwretrieve` program.

If you want save sets to appear separately during retrieve, archive them separately.

Cloned Archives Do Not Appear in Solstice Backup Retrieve Program

When you search for an annotation in the `nwretrieve` program, the Archives attribute does not display archive clones.

To locate the clones, start the query without specifying a Search Annotation attribute. If that query returns too many archives, you can use `mminfo` to locate the archive clone with the same save set ID (ssid) as the archive you want.

Wrong Archive Pool Is Selected

If you create multiple archive pools, the one selected for archive is not the default archive pool. When you create multiple archive pools, the last one created is the one selected for archive.

Second Archive Request Does Not Execute

If you create two archive requests with the same name, only the first request is executed. To avoid this problem, do not create two archive requests with the same name; the newer one will never be executed.

Command-Line Archive Does Not Start Immediately

If you run `nsrarchive` from the command line, the archive does not start immediately after you type the annotation and then [Ctrl]+[D] to start the archive. Wait a short time; there is a delay before the archive starts. Do not press [Ctrl]+[D] multiple times.

Storage Nodes

The following sections describe errors associated with storage nodes.

Storage Node Affinity Errors

If a backup fails and the following message appears:

```
no matching devices; check storage nodes, devices or pools
```

the problem may be related to storage node affinity.

Here are some possible reasons:

- No devices are enabled on the storage nodes in the Storage Nodes list.
- The devices do not have volumes that match the pool required by the backup request.
- All devices are set to read-only.

You must fix the problem and restart the backup. To fix the problem in this example, do one of the following:

- Enable devices on one of the storage nodes on the client's list.
- Correct the pool restrictions for the devices on the storage node list.
- Add another storage node to the list that has enabled devices meeting the pool restrictions.
- Set one of the devices to read/write.
- Adjust the Save Mount Timeout and Save Lockout attributes for the storage node's Device resource. For more information, refer to the online help.

Storage Node Timeout Errors

If `nsrd` initializes on the server and detects that a setting for NSR_MMDCONTROL exists, you will receive the following message:

```
NSR_MMDCONTROL env variable is being ignored

use nsrmmd control timeout attribute instead
```

If you receive this message:

1. Shut down the Solstice Backup daemons.

2. Remove the environment setting for NSR_MMDCONTROL.

3. Restart the Solstice Backup daemons.

4. Start the Solstice Backup Administrator program.

5. Adjust the value of Nsrmmd Control Timeout to the setting that was previously assigned to the NSR_MMDCONTROL variable, or one that best meets your current requirements.

Devices and Autochangers

This section explains how to resolve problems encountered with the use of an autochanger with Solstice Backup software.

Additional Attributes in the Jukebox Resource

The Jukebox resource now contains attributes that provide a detailed view of options that the nsrjb program uses. These are hidden attributes; select View>Details from the Jukeboxes window to display them.

The additional attributes are:

TABLE A-2   Additional Jukebox Attributes

Attribute	Description
Nsrjb Release	Indicates the version of the nsrjb program that is installed on the computer that uses the selected Jukebox resource. The version displayed for Nsrjb Release might differ from the version displayed for the Solstice Backup software on the server, depending on whether you updated your storage nodes to the same Solstice Backup release as the server.
Jukebox Features	Displays the features that your autochanger supports for use by the nsrjb program. If you are installing a software update, the values that were once assigned to the Jukebox Options attribute migrate to the equivalent selections for the new Jukebox Features attribute.
Eject Sleep	Defines the number of seconds an autochanger should remain idle ("sleep") after an eject operation is completed.
Cleaning Delay	Defines the number of seconds an autochanger should remain idle before attempting to unload a cleaning cartridge.
Unload Sleep	Defines the number of seconds an autochanger should remain idle after an unload operation is completed.
Load Sleep	Defines the number of seconds a jukebox should remain idle after a load operation is completed.
Deposit Timeout	Defines the number of seconds a jukebox should wait for a deposit in the mail slot before the nsrjb program abandons further attempts to complete the deposit operation.

Caution – The timed settings are set to the optimal settings for optimal performance. Do not change these settings unless advised to do so by a Technical Support representative.

Maintenance Commands

Solstice Backup device driver software provides the following maintenance commands for diagnosing problems on tape devices and autochangers. For more detailed information regarding these commands, refer to their corresponding man pages.

TABLE A-3 Autochanger Maintenance Commands *(1 of 2)*

Command	Description
lusbinfo	Prints out SCSI information
lusdebug	Sets the library debugging level
lrescan	Rescans for devices
lreset	Resets the SCSI bus
changers	Lists the SCSI autochangers attached to the system
hpflip	Flips the device type of HP optical disk drives
ielem	Initializes the element status
inquire	Lists the devices available
ldunld	Loads or unloads a tape device
msense	Retrieves mode sense data
pmode	Prints mode sense data
relem	Reads the element status
tur	Tests whether the unit is ready
writebuf	Writes a device buffer
sjiielm	Tests the standard jukebox interface (SJI) SJIIELEM command
sjiinq	Tests the SJI SJIINQ command
sjirdp	Tests the SJI SJIRDP command

TABLE A-3 Autochanger Maintenance Commands *(2 of 2)*

Command	Description
`sjirdtag`	Tests the SJI SJIRTAG command
`sjirelem`	Tests the SJI SJIRELEM command
`sjirjc`	Tests the jukebox

▼ To Test the Device Driver Installation

After you install the Solstice Backup device driver software, use the `lusdebug` program to verify the server connection and the `jbexercise` program to test the autochanger. Use the value of the control port assigned to your autochanger (for example, `scsidev@0.6.0`) for *control_port* in the following commands:

```
# lusdebug control_port 0

# jbexercise -c control_port -m model
```

If these commands fail or if you receive error messages, see the following sections for information on the possible cause and solution.

The lusdebug Command Fails

If the `lusdebug` command fails, review these suggestions to identify the potential problems and their solutions:

■ Issue the `sjiinq` command as root, and provide the `control_port` as an argument. You should receive a message similar to the following:

```
scsidev@0.6.0:<EXABYTE EXB-10i EXB-10i >
```

Verify that the information supplied by the message is correct.

If the vendor and model names are incorrect, you supplied the wrong SCSI ID as the device ID during the driver installation. The installation script asks for the SCSI ID of the robot mechanism, not the tape drive.

Uninstall the device driver and reinstall it, and supply the correct address for the autochanger (robotic arm). Make sure that each device on the SCSI bus has a different SCSI ID address.

■ Inspect the following items to verify that the autochanger is properly connected:

a. Make sure all the connectors on the SCSI bus are firmly connected.

b. Make sure none of the SCSI cables are defective.

c. Verify that the SCSI bus is properly terminated and is within the length specified by ANSI SCSI-II specifications (ANSI X3.131-1994).

    Both ends of the SCSI bus must be terminated with the appropriate resistors to be properly terminated. Single-ended SCSI buses are 220 ohms to +5 VDC, 330 ohms to ground. Differential terminators have a 122-ohm characteristic impedance (-5 VDC to +5 VDC). The ends of the SCSI bus are considered to be the last SCSI device at either end of the bus, where both peripheral devices and systems are considered as peer SCSI devices.

    Additional termination (termination placed on devices not at either end of the SCSI bus) is ill-advised. Additional termination causes the hardware bus drivers on each device on the bus to work harder (for example, out of the range of their nominal specification) to affect signal transitions. As a result, they might not be able to meet the timing requirements for some signal transitions.

d. SCSI bus length limitations affect the quality of the signals, thus increasing the likelihood of transmission errors on the bus. For single-ended SCSI buses (the most prevalent), the length is 6 meters, unless Fast SCSI devices are attached and in use, in which case the length limit is 3 meters. This length includes the length of the bus as it is within a device as well as the length of external cables. A reasonable rule of thumb for internal length is to assume 1 meter of internal bus length for the computer chassis and about 0.25 meters per device for external peripheral boxes.

    Differential option SCSI buses can be much longer (due to the electrical differences from single-ended). Allow for a maximum of 25 meters. Never mix differential and single-ended devices.

- Check to see whether an old autochanger driver is still installed. This can be the autochanger driver shipped with earlier versions of Solstice Backup software, or release 1.1 or earlier of the Parity driver, which only supported SCSI bus 0.

  Uninstall the driver according to the instructions shipped with the old driver, and then reinstall the latest version. Special instructions on how to remove the autochanger driver are available in Legato Technical Bulletin 086.

- Check the SCSI IDs on all devices attached to the same bus; make sure that none are the same. If two devices have the same target ID, you can see the following symptoms: SCSI bus reset errors appear in system log files, the machine does not boot, and the `probe-scsi` boot prompt command on SPARC systems hangs.

- If the sensor that verifies whether the tape drive door is open is out of place, follow the instructions provided with your autochanger hardware to determine the problem, or contact your hardware vendor.

- If the autochanger is in sequential mode, change the setting to random mode.

If none of these suggestions resolves the problem, contact Sun Technical Support. Provide the information described in "Information to Gather Before You Call Technical Support" on page 469 and the captured output of the jbexercise, sjiinq, and sjirjc programs. For more information, refer to the *Legato Command Reference Guide* or the jbexercise, sjiinq, and sjirjc man pages.

The jbexercise Command Fails

If the jbexercise command fails, review the following list of suggestions to identify potential problems and their solutions:

■ The jbexercise program prompts you for a nonrewinding device name (for example, on Solaris, /dev/rmt/0mbn). Verify that you have supplied the correct device pathname for the tape drive. The device name must belong to a tape drive in the autochanger, not the autochanger itself.

If you receive the following error message, you did not enter a nonrewinding device name:

```
device not ready
```

■ Make sure that the tape drive for which you enter the pathname works. Insert a volume into the drive and perform the following tests:

a. Use the tar command to copy a small file to the volume.

b. Verify more extensive operations by issuing the tapeexercise command.

If these tests fail, the tape drive is not functioning. For further information on how to configure your tape drive to work with your system, contact your hardware vendor.

If none of these suggestions resolves the problem, contact Sun Technical Support. Provide the information described in "Information to Gather Before You Call Technical Support" on page 469 and the captured output of the jbexercise, sjiinq, and sjirjc programs. For more information, refer to the *Legato Command Reference Guide* or the jbexercise, sjiinq, and sjirjc man pages.

Auto Detected SCSI Jukebox Option Causes Server to Hang

If you install an autodetected SCSI jukebox using jbconfig, and the server hangs:

1. Select the jbconfig option that installs an SJI jukebox. A list of jukeboxes appears.

2. Enter the number that corresponds to the type of jukebox you are installing.

3. Proceed with `jbconfig` until you receive the following message:

```
Jukebox has been added successfully.
```

X11 Libraries Required for `jbconfig`

To use the `jbconfig` program on UNIX, the X11 graphical libraries must be installed. If you attempt to run the `jbconfig` program on a computer that does not have the X11 graphical libraries installed, the following error message is displayed:

```
Symbol S...... in use libvgalaxy.a is not defined
Cannot load libvgalaxy.a
System error, can not run a file.
```

Before you can recover data using an autochanger, you must add and configure the autochanger using the `jbconfig` command.

On an HP-UX computer, ensure that the `/nsr/jukebox` directory and its contents are present before attempting to configure an autochanger. If the directory is not present or it is empty, copy the contents of `/opt/networker/lib` to `/nsr/jukebox` before attempting to configure an autochanger.

Autochanger Inventory Problems

Any of the following situations can cause the autochanger inventory to become outdated:

- The media is manually ejected from the autochanger drive.
- The media is removed from the autochanger.
- The autochanger door is opened.

An outdated inventory means that Solstice Backup software cannot use the autochanger.

To make the autochanger usable again:

1. Verify that the media cartridge is correctly installed in the autochanger and that the autochanger door is closed.

2. Log in as root on the Solstice Backup server.

3. Reset the autochanger by entering the following command at the prompt:

```
# nsrjb -Hv
```

4. Perform an inventory by entering the following command at the prompt:

```
# nsrjb -Iv
```

After the inventory operation is finished, the Solstice Backup server can once again use the autochanger.

For complete information on the use of the nsrjb command, refer to the nsrjb(8) man page or see Chapter 10.

Destination Component Full Messages

The message `Destination component full` usually is the result of a manual operation performed on the autochanger, for example, physically unloading the tape drive by means of the buttons on the autochanger rather than using the Solstice Backup server to unmount the volume. This operation causes the server to lose track of the status of the media in the autochanger.

To resolve the problem, use the Solstice Backup command nsrjb -H to reset the autochanger.

Tapes Are Not Filled to Capacity

You might encounter situations where the Solstice Backup software does not fill tapes to capacity. For example, a tape with an advertised capacity of 4,000 MB can be marked full by the Solstice Backup server after only 3,000 MB of data have been written to it.

To enable the server to use the tape capacity to its fullest, select the highest density device driver appropriate for your device. When a tape is labeled, the server writes to it at the highest density supported by your device.

There are several reasons why the server appears to fill tapes prematurely:

■ Write errors occur during a backup.

Most tape drives try to read after a write operation to verify that the tape was written correctly, and retry if it was not. A write error indicates either an end-of-tape or read error. At any tape error, the Solstice Backup server marks the tape full.

To prevent tape write errors, clean your tape drive regularly and use only data-quality tapes. If cleaning the drive does not seem to help, make sure:

– The device driver is properly configured

– Any necessary switch settings on the tape drive are set to the manufacturer's specifications

– All cabling is secure

– Other potential SCSI problems have been addressed

■ Solstice Backup filemarks take up space on the tape.

The Solstice Backup server periodically writes filemarks to facilitate rapid recovery of data. These filemarks consume varying amounts of tape depending on the type of tape drive. On some drives, filemarks can consume several MB. The number of filemarks the server writes to tape is a function of how many save sets are on the tape. Many small save sets require more filemarks than a few larger ones.

■ Tape capacities vary from tape to tape.

Tape capacities are not constant from tape to tape. Two apparently identical tapes from the same vendor can vary significantly in capacity. This can cause problems if you copy one full tape to another, especially if the destination tape holds less data than the source tape.

■ Data compression affects the tape capacity.

If you use compression on your tape drive, you cannot predict the effect on tape capacity. A compressing drive can provide twice the capacity of a noncompressing drive. The capacity could vary depending on the kind of data being backed up. For example, if a noncompressing drive writes 2 GB of data to a specific tape, the compressing drive could write 10 GB, 2 GB, 5 GB, or some other unpredictable amount of data.

■ Length of tape.

Be sure to verify tape lengths. A 120-meter DAT tape holds more data than a 90-meter DAT tape, and without examining the printed information on the tape cassette carefully, the two tapes can appear identical.

For more detailed information, refer to Technical Bulletin 176, available in the `bulletins.pdf` file included with your Solstice Backup software distribution.

For Solaris, if your tape devices are not directly supported by Sun Microsystems, you will need to re-create your entries in the *st.conf* file. If you need assistance with this, contact Sun Technical Support.

Server Cannot Access Autochanger Control Port

The control port controls the autochanger loading mechanism. Your autochanger's hardware installation manual should have instructions on how to verify whether the control port is properly connected. If you cannot determine whether the control port is working, contact the autochanger vendor for assistance.

Nonrewinding Device Requirement

You must use a nonrewinding device for Solstice Backup backups. The Solstice Backup server writes a filemark on a volume at the end of each backup. When the next backup occurs, the server appends data to the volume based on the position of the filemark. If the device automatically rewinds the data, the file mark position is lost and the data is overwritten by the next backup. You will be unable to recover the previous backup data at a later date.

ATL 2640 Libraries

If a computer connected to an ATL 2640 library is upgraded from a release earlier than 5.5 of Solstice Backup software, perform the following procedure to allow the ATL 2640 library to be used:

1. Set the *polled_import* and *stacker_export* jukebox attributes using the Solstice Backup Administrator program.

   a. Select Media>Jukeboxes to open the Jukeboxes window.

   b. Select View>Details to display hidden attributes.

   c. Select the *Polled_import* and *Stacker_export* options in the Jukebox Features attribute.

   d. Click Apply.

2. Configure the ATL library to enable User Op IOD (inport/outport device). This must be done to allow the IOD door to open to insert tapes.

   a. Connect a DOS/Windows computer to the library DIAG (diagnostic) serial port.

   b. Load the ATL diagnostic program *tapelib.exe* on the computer.

   c. Select Config>User Op IOD, and then select Enable.

In addition, verify that the library is configured with an IOD on the correct side of the device (it must match the physical configuration of the library). Also verify that the pass-through is enabled.

Note – Do not specify any port numbers on the nsrjb command line. In addition, starting with Solstice Backup release 5.5, the numbering scheme for the import/export ports changed from zero-based (0,1,2,3) to 1-based (1,2,3,4) to match the scheme used for all other items (slots and drives).

Depositing Volumes to Slots

Instead of performing an inventory of the autochanger after depositing volumes, nsrjb enables you to specify a range of volume names already contained within the media database.

To deposit a labeled volume that is a member of the Solstice Backup media database into the autochanger, specify only the names of the volumes to be deposited and not the associated port or slot number. The Solstice Backup server will deposit volumes starting from the first port to the first slot that is not empty, for example:

```
# nsrjb volume_name1 volume_name2 volume_name3
```

If you do not want to perform an inventory after depositing labeled volumes to a jukebox that are also members of the Solstice Backup media database, make sure to:

- Specify a range of empty slots for the -S *slots* option. The range specified should begin with the first volume in the range of ports.
- Specify a range of nonempty ports for the -P *ports* option, starting from the first port. If the range of volumes to be deposited does not start from the first port, then the range of ports must be specified using the -P option. For example:

```
# nsrjb -S slots -P ports volume_names
```

Depositing Volumes to Slots on a Solaris Server or Storage Node

When you use the nsrjb -S command to deposit a range of volumes to a range of slots in an autochanger attached to a Solaris server or storage node, the deposit does not occur if the volume names are also specified with the command. To avoid the problem, do not specify the volume names at the command line.

Using the `-o` Command Option for `nsrjb`

If you use the `-o` option twice in the same `nsrjb` command line, the first `-o` option stated will be ignored. For example, in the following example, the `-o notreadonly` task will not be executed:

```
# nsrjb -o notreadonly -o notfull -S 2
```

To perform multiple `nsrjb` tasks with the `-o` option, enter the commands separately. For example, the commands in the previous example must be entered as follows:

```
# nsrjb -Y -o notreadonly -S n

# nsrjb -Y -o notfull -S n
```

Environment Variables Needed for EMASS/Grau and StorageTek Silos

For an EMASS/Grau and StorageTek silos, edit your startup files for Solaris, HP-UX, and AIX. TABLE A-4 describes where the startup files are located.

TABLE A-4 Startup Files

Operating System	Startup File Location
Solaris	/etc/init.d/networker
HP-UX	/sbin/init.d/networker
AIX	/etc/rc.nsr

Locate the following lines in the startup file for your operating system:

```
'start')
(echo 'starting Solstice Backup daemons:') > /dev/console
```

Enter the lines appropriate for your particular silo, as shown in .

TABLE A-5   EMASS/Grau and StorageTek Silo Environment Variables

Silo Model	Lines to Enter
EMASS/Grau	**DAS_SERVER =** *name_of_DAS_server* **export DAS_SERVER** **DAS_CLIENT =** *name_of_system_as_defined_to_DAS_server* **export DAS_CLIENT** **DAS_MEDIUM =** *type_of_tape_drive_used* **export DAS_MEDIUM** **ACI_MEDIA_TYPE =** *type_of_tape_drive_used* **export ACI_MEDIA_TYPE** **Note:** For DAS_MEDIUM and ACI_MEDIA_TYPE, use one of the following values: • 3480 • VHS • OD-Thick • 3590 • OD-Thin • CD • DECDLT • TRAVAN • 8MM • DTF • 4MM • BETACAM • D2 • AUDIOTAPE
StorageTek	**CSI_HOSTNAME =** *name_of_ACSLS_system* **export CSI_HOSTNAME** */solstice_backup_binaries_path/***mini_el &** */solstice_backup_binaries_path/***ssi &**

Diagnostic Tools

A variety of diagnostic tools are available as operating system services and as part of the Solstice Backup product.

To verify that communications sessions can be established, test with ping and rpcinfo, which are tools provided with the operating system software.

Because the Solstice Backup server relies heavily on mapping of ports, use rpcinfo to test the operation of the portmapper. Using ping tests the connection up to the network layer in the OSI model; rpcinfo checks for communication up to the session layer. For instructions on using ping and rpcinfo, see "Client/Server Communications" on page 484.

Contact Sun Technical Support for more tools on testing communications.

Solstice Backup Locale and Code Set Support

Solstice Backup software does not support locales (defined by your operating system) or code sets that remap characters having special meaning for filesystems. Depending on the filesystem, these special characters may include the forwardslash (/), the backslash(\), the colon (:), or the period(.). De_DE.646 is an example of one unsupported locale.

Solstice Backup software might function normally in such an environment, but might not function normally if the locale is changed (the previously existing indexes can become invalid).

The following sections discuss Solstice Backup support of locales and code sets:

- "Locale Support" on page 508
- "Code Set Support" on page 509

Locale Support

Solstice Backup software does not support mixed locale environments, except for the following:

- An English server and a supported localized client
- An English client with a supported localized server

Solstice Backup functionality varies depending on whether the software is localized:

- A localized version of Solstice Backup software on a localized operating system running a supported locale will display localized filenames correctly.
- An English version of Solstice Backup software, running on either a localized or English version of an operating system, might not display non-English filenames correctly.

Code Set Support

TABLE A-6 indicates the Solaris Language environment settings for each supported code set supported in Solstice Backup releases 6.0.1L and later.

TABLE A-6   Solaris Language Environment Settings for Supported Code Sets

Locale	Solaris Language Environment Settings
French	fr fr_FR.ISO8859-1 fr_FR.ISO8859-15
Japanese EUC	ja
Simplified Chinese	zh Chinese

UNIX Platform-Specific Notes

This appendix addresses Solstice Backup software features and functionality specific to one or more UNIX platforms. The main sections in this chapter are:

Solaris

This section provides information specific to Solstice Backup software running on the Solaris platform.

Solstice Backup Executables Not Found in Expected Location for Solaris Client

On Solaris, Solstice Backup executables are installed by default in `/usr/sbin/nsr`. If you start a group backup on a Solstice Backup server that does not have `/usr/sbin/nsr` in the search path for root, the backup fails on a client that has its Solstice Backup executables in `/usr/sbin/nsr`. This is because the `savefs` command is not in the search path.

The best solution is to set the hidden attribute Executable Path for a client that has this problem. To set the Executable Path, display the Clients attribute in detail view and enter the path of the executables, `/usr/sbin/nsr`, in the Executable Path attribute.

Another solution is to modify the search path for root on the Solstice Backup server to include `/usr/sbin/nsr` even if it does not exist locally.

HP-UX

This section provides information specific to Solstice Backup storage node software running on the HP-UX platform.

Autochanger Installation on an HP-UX System

The following sections explain how to install and configure the drivers supplied by Hewlett-Packard.

Selecting SCSI Addresses for the Autochanger

Determine which SCSI address is assigned to each SCSI bus and select SCSI addresses to be allocated to the autochanger drives and controller. To select unused SCSI addresses for your autochanger, log in as root on the Solstice Backup server or storage node and enter the `ioscan -f` command.

SCSI addresses range from 0 to 7. However, SCSI address 7 is reserved for the SCSI interface board in your computer, so you cannot use it. The primary hard disk is usually on SCSI address 6.

<div>⚠️</div>

Caution – Some devices, such as the HP Model 48AL autochanger, choose one SCSI address for the entire autochanger. The 48AL uses a different SCSI logical unit number (LUN) for the device (LUN 0) and robotics (LUN 1). The SCSI LUN appears as the last digit of the H/W Path field in the `ioscan` output.

The following sections provide examples of the command and output to use with different combinations of hardware and operating systems. Follow the example that matches your configuration.

Installing the SCSI Pass-Through Driver

There are two procedures for installing a SCSI pass-through (spt) driver, each depending on the type of hardware you are installing:

- If you are installing a SCSI pass-through driver on an HP 9000 Series 800 containing an HP-PB interface card, see, .

- If you are installing a SCSI pass-through driver on an HP 9000 Series 700 or 800 containing an HSC or PCI interface card, see, .

▼ *To Install an HP-PB or NIO Pass-Through Driver*

For further information regarding the installation of an HP-PB or NIO Pass-Through Driver, refer to the HP-UX man page for scsi_pt.

To install the HP-PB or NIO pass-through driver:

1. Enter the following command to get more information about the peripherals attached to the system you are using:

```
# ioscan -f
```

A screen output similar to the following should appear:

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
ext_bus	7	4/8	c720	CLAIMED	INTERFACE	GSC add-on FAST/WIDE SCSI
target	4	4/8.1	tgt	CLAIMED	DEVICE	
tape	6	4/8.1.0	stape	CLAIMED	DEVICE	STK 4890
ext_bus	8	8/4/4	scsi3	CLAIMED	INTERFACE	HP 28696A-wide SCSE ID=7
target	5	8/4/4.0	target	CLAIMED	DEVICE	
tape	7	8/4/4.0.0	tape2	CLAIMED	DEVICE	Quantum DLT4000
ext_bus	12	10/4/4	scsi3	CLAIMED	INTERFACE	HP 28696A-Wide SCSI ID=7
target	6	10/4/4.6	target	CLAIMED	DEVICE	
unknown	-1	10/4/4.6.0		UNCLAIMED	UNKNOWN	HP C6280-7000

The screen output shows:

- An HSC fast/wide differential SCSI at instance 7, which is connected to an STK 4890 tape drive at SCSI address 1.

- An HP-PB fast/wide differential SCSI adapter at instance 8, which is connected to a Quantum DLT4000 tape driver at SCSI address 0.

- An HP-PB fast/wide differential SCSI adapter at instance 12, which is connected to the robotics controller for an HP C6280-7000 at address 6.

The robotics controller is listed as UNCLAIMED/UNKNOWN, because the HP-PB pass-through has not been configured to communicate with it.

2. To complete the configuration, go to the build area:

```
# cd /stand/build
```

3. Create a system from the existing kernel:

```
# /usr/lbin/sysadm/system_prep -s system
```

Check whether the SCSI pass-through driver is already configured into the current kernel. The driver is represented by the name *spt*:

```
# grep spt system
```

4. If no *spt* line prints, add a driver statement to the SCSI drivers area of the system file, for example, /stand/build/system:

```
# spt
```

5. Append a driver line to the system file, in this format:

```
driver path spt
```

where *path* is the complete hardware path of the desired device (from the ioscan command).

If you were using the previous ioscan display with the library robotics connected to the hardware path 10/4/4.6.0, the line would read:

```
driver 10/4/4.6.0
```

A driver statement is required because the system cannot autoconfigure the SCSI pass-through driver. One statement is required for the robotics controller on each library.

Note: Adding a line such as driver *path* spt might prevent the system from recognizing tape drives added to that SCSI bus. You might have to add lines for driver *path* spt for each tape drive that your system was previously able to recognize.

6. Save the old system file:

```
# mv /stand/system /stand/system.prev
```

7. Move the new system to the current one:

```
# mv ./system /stand/system
```

8. Build a new kernel:

```
# /usr/sbin/mk_kernel
```

9. Verify that the pass-through driver has been built into the new kernel:

```
# what ./vmunix_test | grep scsi_pt
```

A line similar to the following appears:

```
scsi_pt.c Revision: 10.20 $Date: 97/07/18 14:38:27 $
```

10. Save the old kernel:

```
# mv /stand/vmunix /stand/vmunix.prev
```

11. Move the new kernel to be the current one:

```
# mv /vmunix_test /stand/vmunix
```

12. Boot the system from the new kernel:

```
# /usr/sbin/shutdown -r 0
```

13. Verify the peripheral configuration. After the system reboots, enter the following command:

```
# ioscan -f -C spt
```

Output from the ioscan should resemble the following:

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
spt	0	10/4/4.6.0	SPT	CLAIMED	device	HP C6280-7000

14. Proceed to "To Create a Device File" on page 517.

▼ To Install a GSC, HSC, or PCI Pass-Through Driver

The following procedure assumes you are using the SAM terminal mode.

To install a GSC, HSC, or PCI pass-through driver:

1. Run SAM.

2. Select Kernel Config and press [Return].

3. Select Drivers and press [Return].

4. Select SCTL from the list. The SCSI_ctl driver is represented by the name SCTL.

   If the current state is in, proceed to "To Create a Device File" on page 517.

   Note: You can choose any name for the device as long as it is not something reserved by the system. For example, do not use a name such as /dev/null.

5. From the menu bar, select Actions>Add Drivers to Kernel and press [Return].

6. From the menu bar, select Actions>Create a New Kernel and press [Return].

7. Respond Yes to the prompt "Are you sure?" and press [Return].

8. The Creating Kernel message appears (for approximately 1 to 2 minutes), followed by the Move Kernel Message. Select OK and press [Return]. The system reboots.

9. Proceed to "To Create a Device File" on page 517.

Installing the Autochanger Hardware

To install the autochanger hardware, follow the instructions in the appropriate hardware installation manuals, as provided by your vendor. Use the SCSI addresses that you selected in the earlier section, "Selecting SCSI Addresses for the Autochanger" on page 512, to set any required switches on the autochanger for the devices and robotics controller.

For example, using the previous Series 800 example, you set the robotics SCSI address on the autochanger to 3 and the device SCSI address to 4, according to the instructions in the hardware manual supplied by your vendor.

▼ To Shut Down the System

Before installing the autochanger device hardware, shut down your system.

To shut down the server:

1. Halt the server by entering the following command:

```
# sync

# shutdown -h time
```

   where *time* is the grace period, in seconds, for users to log off before the system shuts down. The default grace period is 60 seconds. If *time* is zero, shutdown gives users very little time to log out but it runs more quickly.

2. Turn the main power switch of the server to the Off position.

▼ To Start the System

To start the system:

1. Turn the main power switch of the autochanger to the On position.

2. Turn the main power switch of the server to the On position. This reboots the server, as required after installing new autochanger hardware.

▼ To Create a Device File

To create a device file:

1. Verify that the *spt* has been successfully installed with the following command:

```
# ioscan -kfn
```

2. Verify that the autochanger has been claimed by the driver. If the driver has been claimed, CLAIMED should appear under the S/W State header. If not, you must verify that the installation has been completed properly.

3. If the device entry is already completed by the operating system, then use the OS-defined entry and proceed to verify the installation.

4. If the device entry is not defined, you must create a device file. Use the `mknod` command to create a device file to access the robotics controller. The command syntax is:

```
/user/sbin/mknod /dev/sji devfilename c majornum minornum
```

where:

- *devfilename* is the user defined name of the device file;
- *majornum* is the character major number from the `lsdev` command;
- *minornum* is the minor number; it must be entered in this format:

```
0xIITL00
```

  - `II` is the two-digit card instance number, in hexadecimal (from `ioscan` command; proper `ext_bus` entry)
  - `T` is the largest SCSI ID number
  - `L` is the LUN number
  - `00` is reserved

Major Number

To determine the value for *majornum*, enter one of the following commands:

- If the robotics is attached to an HP-PB adapter:

```
# lsdev -d spt
```

- If the robotics is attached to an HSC or PCI adapter:

```
# lsdev -d sctl
```

The output should resemble the following. The assigned number may differ from those displayed in this example:

	Character	Block	Driver	Class
HP-PB	75	-1	spt	spt
HSC or PCI	203	-1	sctl	ctl

The value for *majornum* is the number in the Character column.

Minor Number

To determine the value for *minornum*, use the ioscan command. The relevant lines in the ioscan output are those:

- Pertaining to the controller itself (which will contain HP C6280-7000 in the Description column).
- Pertaining to the adapter to which the controller is connected (which will be the second line above the line for the controller and contains "ext_bus" in the class column). For example, if the robotics is attached to an HP-PB adapter card, the ioscan should resemble the following:

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
ext_bus	12	10/4/4	scsi3	CLAIMED	INTERRACE	HP 28696-Wide SCSI ID=7
target	6	10/4/4.0	target	CLAIMED	DEVICE	
spt	0	10/4/4.6.0	spt	CLAIMED	DEVICE	HP C6280-7000

Given the information in the previous example, the correct mknod command would be:

```
# /usr/sbin/mknod /dev/sji dev_filename c 75 0x0c6000
```

In the case of an HSC or PCI adapter, the ioscan output varies from the previous in one of two ways.

If the schgr driver is configured on the system, it appears associated with the library. The `ioscan` output line resembles:

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
spt	0	10/4/4.6.0	schgr	CLAIMED	DEVICE	HP c6280-7000

If the schgr driver is not configured on the system, no driver appears to be associated with the library. The `ioscan` output line resembles:

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
unknown	-1	10/4/4.6.0	schgr	UNCLAIMED	DEVICE	HP c6280-7000

▼ To Test the Device Driver and Device File Installation

After you have installed the device driver and created the device file, run the `inquire` command to list available SCSI devices:

```
# inquire
```

The following is an example of the output from this command:

```
scsidev@0.1.0:HP      C1194F              0.14|Autochanger (Jukebox),
/dev/rac/c0t1d0

scsidev@0.2.0:Quantum DLT4000            CC37|Tape,
/dev/rmt/c0t2d0BESTnb

scsidev@0.3.0:Quantum DLT4000            CC37|Tape,
/dev/rmt/c0t3d0BESTnb

scsidev@0.4.0:Quantum DLT4000            CC37|Tape,
/dev/rmt/c0t4d0BESTnb

scsidev@0.5.0:Quantum DLT4000            CC37|Tape,
/dev/rmt/c0t5d0BESTnb
```

▼ To Test the Autochanger Connection

To test the autochanger connection, run the `jbexercise` program with at least two pieces of "scratch" media loaded in the first and last slots of the autochanger. The devices should be empty, with the doors open. For complete instructions, refer to the `jbexercise` man page.

When running `jbexercise`, specify the control port and the autochanger model type. The control port for SCSI autochanger models is typically `/dev/sjid1u1` (or `/dev/sjid1u2` for a second autochanger). The following command runs the `jbexercise` program on an Exabyte 10i autochanger:

```
# jbexercise -c (/dev/sjid1u1) -m EXABYTE_10I
```

"Hardware Address Not Found" Message Encountered

The HP-UX system configuration logger might generate the error message:

```
bootpd pid# Hardware address not found hardware-address
```

A similar message is written to the Solstice Backup `/nsr/logs/messages` file. If you encounter this message:

1. Log in as root to edit the `/etc/syslog.conf` file and change every instance of daemon.notice to local7.notice.

2. Edit the Notifications resource configured for the Log Default notification and change the value shown in the Action attribute from daemon.notice to local 7.notice.

3. Apply and save the changes to the Log Default notification.

4. Enter the following command at the prompt:

```
# cat /etc/syslog.pid
```

5. Use the pid# obtained from the /etc/syslog.pid file to kill the designated pid# and cause the syslogd daemon to reread its configuration file:

```
# kill -HPU pid#
```

The local7 facility is provided as an example. View the syslog.h system header file to determine other possibilities for the facility. By default, LOG_LOCAL0 to LOG_LOCAL7 are reserved for local use. If they are not used by other local utilities on your system, the Solstice Backup server can use them to avoid the hardware address problems encountered with bootpd.

Errors from Unsupported Media in HP Tape Drives

Certain HP tape drives can only read 4-mm tapes of a specific length. Some, for example, read only 60-meter tapes and do not support the use of 90- or 120- meter tapes. To determine the type of tape your HP drive supports, consult the hardware manual provided with the drive.

If you attempt to use unsupported media in an HP tape drive, you might encounter the following types of error messages in the respective situations:

■ When you use the nsrmm or nsrjb command to label the tape:

```
nsrmm: error, label write, No more processes (5)
```

■ When you attempt to use the scanner -i command:

```
scanner: error, tape label read, No more processes (11)

scanning for valid records …

read: 0 bytes

read: 0 bytes

read: 0 bytes
```

Unloading Tape Drives on an HP-UX Storage Node

When you use the `nsrjb -u -S` command to unload a tape drive in an autochanger attached to an HP-UX 10.x storage node, all of the tape drives inside the autochanger are unloaded to their respective slots. If you only want to unload a single drive to its corresponding slot, use the `nsrjb -u -f` *devicename* command instead.

Note – When you use the `nsrjb` command with the `-f` *devicename* attribute and your device name contains backslashes, you must use the backslash "escape" character with the backslash. Therefore, a device with the name `rd=sn1:\\.\Tape0` would be entered at the command line as `rd=sn1:\\\\.\\Tape0`.

AIX

This section provides information specific to Solstice Backup storage node software running on the AIX platform.

▼ To Activate Full Diagnostic Core Dumps on AIX

Earlier releases of Solstice Backup software on an AIX system do not automatically provide full diagnostic core dumps.

You need to set the NSR_FULLDUMP environment variable to 1 to activate full diagnostic core dumps.

To set the variable and the core file size:

1. Log in as root and set the environment variable with one of the following commands as appropriate for your shell tool:

```
# export NSR_FULLDUMP=1

# setenv NSR_FULLDUMP 1
```

2. Use the Systems Management Interface Tool (SMIT) or edit the `/etc/security/limits` file to set the core file size to -1 for root or default.

3. Set the Enable Full CORE Dump option to True (select Change or Show Characteristics of Operating System from the System Environments choices).

   For the core file size change to take effect, you might need to reboot your system after you edit the limits file. You can check the file size limit with the `ulimits -a` command. For the size that you set in the procedure described, the `ulimits` command should return a value of "unlimited."

STK-9840 Drives Attached to AIX

If you attach an STK-9840 drive to an AIX server for Solstice Backup backups, use SMIT to modify the IBM tape drive definition field to set the value of USE EXTENDED FILE MARK to YES.

IRIX

This section provides information specific to Solstice Backup storage node software running on the IRIX platform.

IRIX Displays "No Space Left in Tape" Message

If you use an IRIX Solstice Backup storage node, you might receive the following message:

```
BSF invalid argument no space left in tape
```

This is not a Solstice Backup error, but an indication that the tape drive you are using might not be configured within the machine's kernel. You need to ensure that the device is supported and configured for your machine by applying a kernel patch, available from the Silicon Graphics web site.

To resolve the problem:

1. Go to the web page `www.sgi.com/support/patch_intro.html`.

2. Download the recommended patch.

3. Follow the instructions provided to apply the patch set.

SGI tests specific firmware revisions for each drive and robot with IRIX Solstice Backup software. Even if your site uses drives and robotics that SGI has qualified, you might need to apply patches to IRIX to make the combination work.

Visit the SGI Services and Support web site for information about and distribution of patches required for your particular configuration. You can also obtain information regarding the current list of supported drives, robotics, and firmware revisions supported by IRIX.

Glossary

This glossary provides definitions for terms used in this guide.

ACL	An abbreviation for access control list. This is a list that specifies the permissions assigned to a specific file or directory.
	To recover a file that has an associated ACL, you must either be logged into the system as root, as Administrator, or as the file's owner.
active group	A Solstice Backup backup group that has its Autostart attribute enabled.
administrator	The person normally responsible for installing, configuring, and maintaining Solstice Backup software.
Administrators group	A Windows NT and Windows 2000 user group whose members have all the rights and abilities of users in other groups, plus the ability to create and manage all the users and groups in the domain. Only members of the Administrators group can modify Windows NT and Windows 2000 operating system files, maintain the built-in groups, and grant additional rights to groups.
annotation	A comment that you associate with an archive save set to help identify that data later. Annotations are stored in the media index for ease of searching and are limited to 1,024 characters.
archive	The process by which Solstice Backup software backs up directories or files to an archive volume and then grooms them to free disk space. When data is archived, it is written to one or more storage volumes and then marked so that it is never subject to automatic recycling. You can delete the archived files from the client, thus freeing disk space. See also *grooming*.
archive clone pool	A pool composed exclusively of archive clone save sets.

archive pool	A volume *pool* composed exclusively of archive save sets. Archived save sets are in a different format than regular backup save sets, and must be maintained on separate media.
archive volume	A tape or other storage medium used to store Solstice Backup archive data, as opposed to a *backup volume*.
ASM	An abbreviation for application-specific module. An ASM is a program that, when used in a directive, specifies the way that a set of files or directories is to be backed up and recovered. For example, **compressasm** is a Solstice Backup directive used to compress and decompress files.
ASM specification	A directive that specifies how files or directories with a matching pattern are backed up. This specification appears in the form [+] ASM [args ...]: pattern ... For more information, refer to the **nsr_5** man page or the *Sun Command Reference Guide*.
attribute	A feature of a resource. It is a service or information that the *resource* provides.
authorization code	A code that is unique to your network that unlocks the software for permanent use.
autochanger	A mechanism that uses a robotic arm to move media among various components located in a device, including slots, media drives, media access ports, and transports. Autochangers automate media loading and mounting functions during backup and recovery. The term autochanger refers to a variety of robotic libraries, including autoloader, *carousel*, datawheel, jukebox, library, and near-line storage.
auto media management	A feature that enables the storage device controlled by the Solstice Backup server to automatically label, mount, and overwrite a volume it considers unlabeled. Volumes that are eligible for reuse are also automatically recycled.
backup	The writing of saved data to a volume.
backup cycle	The period of time from one level full backup to the next level full backup.
backup group	See *group*.
backup level	See *level*.
Backup Operators group	A group of Windows NT or Windows 2000 users who can log on to a domain from a computer or a server, back it up, and restore the data. Backup operators can also shut down servers or computers.

backup volume	A tape or other storage medium used to store Solstice Backup backup data, as opposed to an archive volume.
base enabler code	See *enabler code*.
bootstrap	A save set that is essential for the Solstice Backup disaster recovery procedures. The bootstrap is composed of three components that reside on the Solstice Backup server: the *media database*, the *resource database*, and a *server index*.
browse policy	A policy that determines how long entries for your backup data remain in the client file index.
carousel	A tray or tape cartridge that holds multiple backup volumes.
client	A computer that accesses the Solstice Backup server to back up or recover files. Clients may be workstations, computers, or fileservers.
client file index	A database of information maintained by the Solstice Backup server that tracks every database object, file, or *filesystem* backed up. The Solstice Backup server maintains a single client index file for each client computer.
client-initiated backup	See *manual backup*.
clone	The Solstice Backup process used to make an exact copy of saved data (save sets). You can clone individual save sets or the entire contents of a backup volume. Cloning is different from a simple copy operation carried out on an operating system or hardware device. It is indexed and tracked by Solstice Backup software in both the *client file index* and the *media database*.
clone pool	A pool of volumes composed exclusively of cloned data. Three types of clone pools can be used: backup clone, archive clone, and migration clone. Save sets of different types (for example, archive and migration) cannot be intermixed on the same clone volume.
clone volume	A volume belonging to a clone pool.
connection port	The port Solstice Backup processes use to perform backup and recovery sessions through a firewall.
consolidate	To create a complete backup of a save set by merging the most recent level 1 save set with its corresponding full level save set. For more information, refer to the **nsrssc** man page or the *Sun Command Reference Guide*.

continued save set	Data associated with a save set that is continued from a previous volume. Continued save sets are created by the backup server when large save sets are being backed up, cloned, or archived to multiple volumes.
daemon	A program that lies dormant waiting for a specified condition to occur.
datawheel	See *autochanger*.
device	1. A storage unit that reads from and writes to storage volumes (see volume). A storage unit can be a tape device, optical drive, *autochanger*, or file connected to the *server* or *storage node*. 2. When Dynamic Drive Sharing (DDS) is enabled, refers to the access path to the physical drive.
DFS component	An abbreviation for Distributed File System root or child node. DFS is a Microsoft add-on for Windows NT 4.0 Server (Alpha or Intel) or Windows 2000 that allows you to create a logical directory of shared directories that span multiple machines across a network.
directed recover	A method of recovery that recovers data that originated on one client computer and re-creates it on another client computer.
directive	An instruction that directs Solstice Backup software to take special actions on a given set of files for a specified client during a backup.
drive	When Dynamic Drive Sharing (DDS) is enabled, refers to the physical backup object, such as a tape drive, disk, or file. See also *device*.
enabler code	A special code provided by Sun that activates the software. The enabler code that unlocks the base features for software you purchase is referred to as a base enabler. Enabler codes for additional features or products (for example, autochanger support) are referred to as add-on enablers.
exit code	An indicator that specifies whether a backup or recovery session succeeded. An exit code of zero ("0") indicates the session completed successfully. A nonzero exit code indicates the session did not complete successfully.
expiration date	The date when the volume changes from "read/write" to "read-only."
expired save set	A save set whose browse time has been reached; therefore, it can no longer be browsed. In addition, the save set has been removed from the client file index.

file index	See *client file index*.
filesystem	1. A file tree that is on a specific disk partition or other mount point. 2. The entire set of all files. 3. A method of storing files.
firewall	A system designed to prevent unauthorized access to or from a private network. All messages entering or leaving the intranet pass through the firewall, which examines each message and blocks those that do not meet the specified security criteria. There are several types of firewall techniques. The Solstice Backup software supports client backups from computers that are protected by *packet filtering*.
fork	A child process created by Solstice Backup software to perform the requested operation. In instances where a command uses a parallelism value, Solstice Backup software creates multiple instances of that command. Each instance of the command is identical to the original command and is referred to as a child process. Once each child process is created, they are run simultaneously.
full backup	See *level*.
grooming	The Solstice Backup process of removing the original files from a local disk after a successful archive operation.
group	A client or group of clients configured to start backing up files to the Solstice Backup server at a designated time of day.
inactivity timeout	An attribute that indicates the number of minutes Solstice Backup software waits before determining that a client is unavailable for backup.
incremental	See *level*.
index policy	A policy that specifies how long file and data index entries should remain valid. Clients specify which index policies they wish to use.
jukebox	See *autochanger*.
LCP	An abbreviation for library control program, a program that the Legato GEMS SmartMedia software uses to control robotic libraries.

level	A measurement that determines how much data Solstice Backup software saves during a scheduled or manual backup.
	A full (f) backup backs up all files, regardless of whether they have changed. Levels one through nine [1-9] back up files that have changed since the last lower numbered backup level. An incremental (incr) backup backs up only files that have changed since the last backup.
library	See *autochanger*.
license enabler	The enabler code that enables you to run a feature or product.
local host	The node on which the client or server program is running.
LUS	An abbreviation for Legato User SCSI. The LUS driver is used by Sun software products as a proprietary device driver that sends arbitrary SCSI commands to an autochanger.
manual backup	A backup that a user requests from the client's **save** program. The user specifies participating files, filesystems, and directories. A manual backup does not generate a *bootstrap* save set.
media	The physical storage medium to which backup data is written. Solstice Backup software supports tape, magnetic or optical disk, and filesystems as backup media. See also *volume*.
media database	A database that contains indexed entries about the storage volume location and the life cycle status of all data and volumes managed by the Solstice Backup server. See also *volume*.
migration	The process of moving data from a local filesystem to storage media in the migration store to free up disk space on the local drive.
multiplexing	A Solstice Backup feature that permits data from more than one save set to be simultaneously written to the same storage device.
NDMP	An abbreviation for Network Data Management Protocol, which is a storage management client/server protocol for enterprise-wide backup of network-attached storage. Solstice Backup software uses NDMP to provide connections to computers with NDMP data modules for tape operations, allowing a significant reduction in network traffic.
near-line storage	See *autochanger*.
Solstice Backup client	See *client*.

Solstice Backup server	See *server*.
Solstice Backup storage node	See *storage node*.
NFS client	A computer that can access files on a network file system (NFS) server.
NFS server	A computer that contains exported filesystems that NFS clients can access.
nonclone pool	Pools that contain data that has not been cloned.
notification	A message generated and sent to the Solstice Backup administrator about important Solstice Backup events.
online indexes	The databases located on the Solstice Backup server that contain all the information pertaining to the client backups (*client file index*) and backup volumes (*media database*).
operator	The person who monitors the server status, loads backup volumes into the server devices, and otherwise executes the day-to-day Solstice Backup tasks.
override	A Solstice Backup feature that allows you to configure a different backup level for a specific date listed in a Schedule resource.
packet filtering	A method of firewall protection that looks at each packet entering or leaving the network and accepts or rejects it based on user-defined rules. See also *firewall*.
parallelism	A Solstice Backup feature that enables the Solstice Backup server to either back up save sets from several clients or many save sets from one client at the same time. Parallelism is also available during recovers.
pathname	A set of instructions to the operating system for accessing a file. An *absolute pathname* tells how to find a file beginning at the root directory and working down the directory tree. A *relative pathname* tells how to find the file starting where you are now.
policy	A set of constraints that specify how long an entry can remain in a client's online file index. When a policy expires, the save sets associated with that policy are marked recyclable. Each client resource uses two policies: a browse policy and a retention policy. For more information, refer to the **nsr_policy (5)** man page or the *Sun Command Reference Guide*.

pool	A feature that enables you to sort backup data to selected volumes. A pool contains a collection of backup volumes to which specific data has been backed up.
probe	The process Solstice Backup software uses to determine the directories or files to back up on each client.
purging	The process of deleting all entries for files on the volume from the client file index, but allowing entries for the save sets to remain in the media database.
recover	A recovery method that re-creates an image of the client filesystems and database on the Solstice Backup server.
recyclable save set	A save set whose browse and retention policies have been reached; therefore, the save set has been removed from the media database.
recyclable volume	A volume whose data has passed both its browse and retention policies and is now available for relabeling and use by a Solstice Backup server or storage node.
Registry	A database of configuration information central to Windows NT and Windows 2000 operations. It centralizes all Windows settings and provides security and control over system, security, and user account settings.
remote device	A storage device that is attached to a Solstice Backup storage node.
RPC	An abbreviation for remote procedure call, which is the protocol the Solstice Backup server uses to perform client requests over a network.
resource	A component of Solstice Backup software that describes the Solstice Backup server and its clients. Devices, schedules, clients, groups, and policies are examples of Solstice Backup resources. Each resource contains a list of attributes that define the parameters to use for the specific Solstice Backup resource.
resource database	A database that contains information about each of the configured backup server's resources.
retention policy	A policy that determines how long save set entries are retained in the Solstice Backup server's media database.
retrieve	The process of locating and copying back files and directories that Solstice Backup software has archived.

retry mechanism	The action Solstice Backup software performs when client operations fail. This situation might occur when the rate of transmission is either low or nonexistent. By using this mechanism, a previously failed operation might be more successful. Another common situation whereby a retry mechanism might succeed is when the client is in a reboot cycle.
root	1. (UNIX only) The UNIX superuser account (with user name "root" and user ID). By extension, the privileged system-maintenance login on any operating system. 2. (Windows NT/Windows 2000 and UNIX) The top node of the system directory structure; the home directory of the root user.
save set	A group of files or a filesystem from a single client computer backed up onto storage media.
save set consolidation	The process that merges a level 1 backup with the last full backup of a save set to create a new full backup. See also *level*.
save set ID	An internal identification number that Solstice Backup software assigns to a save set.
save set recover	The recovery of specified save sets to the Solstice Backup server.
save set status	The Solstice Backup attribute that indicates whether a given save set is restorable, recoverable, or recyclable. The save set status also indicates whether the save set has been successfully backed up.
save stream	The data and save set information being written to a storage volume during a backup. A save stream originates from a single save set.
server	The computer on a network that runs the Solstice Backup server software, contains the online indexes, and provides backup and recovery services to the clients and storage nodes on the same network.
server index	A file that lists all the server files that were backed up during a scheduled backup.
service port	The port used by a server or storage node to listen for backup and recovery requests from clients through a firewall.
shell prompt	A cue for input in a shell window where you enter a command.
silo	A repository for holding hundreds or thousands of volumes. Silo volumes are identified by barcodes, not by slot numbers.

skip	A backup level in which designated files are not backed up. See also *level*.
SMS	An abbreviation for system management software, which is a Microsoft-based software installation system that allows the wide-scale, automatic installation of software products on clients from a single remote server.
SNMP	An abbreviation for Simple Network Management Protocol, which is a protocol that defines the communication between a manager (sometimes called a monitor or management station) and an object (the item being managed). Solstice Backup software uses SNMP to send messages to the administrator about Solstice Backup events.
ssid	See *save set ID*.
staging	The process of moving data from one storage medium to another, less costly medium, and later removing the data from its original location.
stand-alone device	A *storage device* that contains a single drive for backing up data.
storage device	The hardware that reads and writes data during backup, recovery, or other Solstice Backup operations.
storage node	A storage device physically attached to another computer whose backup operations are administered from the controlling Solstice Backup server.
versions	The date-stamped collection of available backups for any single file.
volume	A unit of storage media, such as a magnetic tape, an optical disk, or a file. A storage device reads from and writes to volumes, which can be physical units (for example, a labeled tape cartridge) or logical units (for example, optical media can store multiple volumes on a single physical platter).
volume ID	The internal identification assigned to a backup volume by Solstice Backup software.
volume name	The name you assign to a backup volume when it is labeled.
volume pool	See *pool*.

Index