Solstice Backup™ 6.1

Release Notes

Please
Recycle

™
Adobe PostScript

Contents

Tables

Preface

The *Solstice Backup 6.1 Release Notes*, contains brief descriptions of new features, known limitations, and configuration tips and workarounds for this release of Solstice Backup software.

- Chapter 1 contains information about the release in general.
- Chapter 2 contains information specifically related to the Network Data Management Protocol (NDMP) module.
- Chapters 3 contains information on installing and configuring the Solstice Backup software in a Sun Cluster 3.0 environment.
- Chapters 4 contains information on administering Solstice Backup software in a Sun Cluster 3.0 environment.
- Appendix A describes non-English versions of Solstice Backup software.
- Appendix B describes how to map logical device paths to physical drives.
- Appendix C describes Solstice Backup 6.1 updating strategies.

This document was originally published in August, 2001.

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

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

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

Related Documentation

Application	Title	Part Number
Administration	*Solstice Backup 6.1 Administrator's Guide*	816-1412-10
Installation Guide	*Solstice Backup 6.1 Installation Guide*	816-1413-11

Accessing Sun Documentation Online

The `docs.sun.com`™ 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 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 (816-1413-11) of your document in the subject line of your e-mail.

Solstice Backup New Features and Improvements

Solstice Backup release 6.1 contains new features and improvements that are described in this chapter. For more information, refer to the *Solstice Backup 6.1 Administrator's Guide* and the *Solstice Backup 6.1 Installation Guide*.

This chapter includes the following topics:

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

Supported Operating Systems

The following operating system platforms are supported in this release:

- IBM AIX® 4.3, 4.31, 4.3.2, 4.3.3 (IBM RS6000/PowerPC), AIX 5L version 5.1
- Sun Solaris™ 2.6, 7, 8
- HP-UX® 10.20, and 11.0, 11i (HP 9000 Series)
- Compaq® Tru64™ UNIX 4.0 (F-G), 5.0A, 5.1
- Linux® (Red Hat™ 6.2, 7.0 [client only], 7.1); Caldera 2.4, 3.1; SuSE® 7.0, 7.1; Turbolinux® 6.0, 6.1(Intel x86 only)
- SGI IRIX 6.5, 6.5.1-12
- Symmetry: DYNIX/ptx 4.4.4, 4.4.6-4.4.8
- NUMA-Q: DYNX/ptx 4.4.4, 4.4.6-4.4.8, 4.5.1

For further up-to-date details on supported operating systems, refer to the Legato Compatibility Guides available on the Legato web site at www.legato.com.

New or Enhanced Features

The following features are new or enhanced in release 6.1 with Patch 111962-xx:

Support for Sun Cluster 3.0

The Solstice Backup software can be used for backups within a Sun Cluster environment. By configuring the Solstice Backup server software as a failover resource group within a Sun Cluster environment, it becomes a higly available application. For the complete installation and configuration instructions, see .

A Solstice Backup server installed within a cluster can migrate or fail over among other nodes in the same cluster. Fail over allows another node in the cluster to take over the operations from the failed node. When fail over occurs during a scheduled backup, the interrupted save operation restarts when the fail over has completed, beginning with the interrupted save set.

Support for LTO Cartridges in an ADIC DAS 3.02 Environment (Legato Bug ID: LGTpa30293)

This release supports Linear Tape Open (LTO) cartridges in an ADIC Dual Attach Station (DAS) 3.02 environment. The Solstice Backup software has been updated to query DAS on the ADIC silo for the required volume media type. Modifications to the media type query method will allow the Solstice Backup software to automatically handle new media types when available.

Additionally, the *DAS_MEDIUM* and the *ACI_MEDIA_TYPE* environment variables have been updated to include all existing media types supported by the Solstice Backup software.

Removal of High Speed Device Licensing Restriction (Legato Bug ID: LGTpa30949)

Licensing enforcement for high speed devices has been removed from the Solstice Backup software.

The following devices have been defined as high speed:

- Sony™ DTF
- Ampex™ DST
- IBM® 3590
- Storage Tek™ SD-3

Bull Fiber Channel and Emulex 7000e/8000 on AIX 4.3.3 (Legato Bug ID: LGTpa28460)

The Solstice Backup software now supports the Emulex 7000e/8000 Host Bus Adapters (HBA) on AIX 4.3.3 when using the Bull Fiber Channel (FC) driver.

The Bull FC implementation created a new device object `/dev/fcp*`. Modifications were made to the Solstice Backup `lus` configuration method to recognize `/dev/fcp*` device objects.

Note – For more information on the Bull driver, refer to the Bull web site at `www.bull.com`

Note – For more information, refer to the SAN Hardware Compatibility Guide on the Legato web site at `www.legato.com`.

IBM Native Fiber Channel Driver Support on AIX (Legato Bug ID: LGTpa31238)

This release supports the IBM Native Fiber Channel (FC) driver providing a certified Storage Area Network (SAN) solution on AIX. Modifications were made to the `lus` configuration method to recognize the IBM Native FC driver.

Note – For more information, refer to the SAN Hardware Compatibility Guides on the Legato website at `www.legato.com`.

NDMP Recovery Fixes Combined and Available for Solstice Backup 6.1 (Legato Bug IDs: LGTpa30748, 31270) (Resolved in Solstice Backup 6.1 with Patch 111962-xx)

The fixes are available in this patch.

Fiber Channel Support for Linux (Legato Bug ID: LGTpa31755)

This release supports Fiber Channel Host Bus Adaptor (HBA) on Linux.

Note – For more information, refer to the SAN Hardware Compatibility Guide on the Legato web site at `www.legato.com`.

Large file Support for Linux 86 Platforms, Kernel Version 2.4 (Legato Bug ID: LGTpa31267)

This release supports files greater than 2GB on Linux 86 platforms, kernel version 2.4. Support for files greater than 2GB was possible in previous releases of the Solstice Backup software; however this functionality previously required various patches.

Localized Versions of the Solstice Backup Software

Solstice Backup 6.1 L10N will be available in Japanese, Chinese and French. The media kit will include all English language packages and the appropriate language packages.

To install and run any version of the non-English, localized Solstice Backup software, you must first install the English language release of the Solstice Backup software. Once this release of the software is installed, you can install the required localized version of the Solstice Backup software. For more information, see

Dynamic Drive Sharing

Dynamic Drive Sharing (DDS) is an optional feature of Solstice Backup 6.1 that allows tape drives to be shared among multiple hosts (such as a server or storage node) within a single Solstice Backup datazone. All hosts sharing drives must be on the same Fiber Channel SAN.

The feature is not supported within a cluster environment.

The DDS option is licensed per drive; each drive that is to be shared requires a license.

For further information about enabling DDS with the Solstice Backup NDMP module, see .

The jbverify Program

The jbverify program ensures that devices defined in the Solstice Backup database are configured properly. It does this by checking for accessibility and usability. By default, the program checks all devices in the Solstice Backup database. The program can also be configured to check only:

- Autochangers (using the -j option)
- Stand-alone drives (using the -d option)
- Local devices (using the -l option)
- Individual jukeboxes (using the -J option)
- Specific drives (using the -D option)

Note – Legato GEMS SmartMedia devices are not tested by the jbverify program. Legato GEMS SmartMedia is currently available only from Legato Systems, Inc. For more information about this product, contact Legato Systems at sun@legato.com.

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.

The jbverify program can be used:

- 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.
- To test devices remotely from a storage node that is not a Solstice Backup server.

Linux Platform Support

In addition to the platforms listed in the *Legato NetWorker Installation Guide, Release 6.1, Linux Version*, the product offering for Linux now includes support for additional platforms:

- Caldera 3.1, kernel version 2.4.2
- Red Hat 7.0 (client only), kernel version 2.2.16
- Red Hat 7.1, kernel version 2.4.2
- SuSE 7.0, kernel version 2.2.16
- SuSE 7.1, kernel version 2.4.2
- Turbolinux 6.0, 6.1, kernel version 2.2.13

Information regarding requirements of other modules or libraries has also been updated. For system requirements for various Linux platforms, refer to the *Legato NetWorker Installation Guide, Release 6.1, Linux Version*.

Fixed Bugs

The following tables describe bugs fixed in Solstice Backup release 6.1.

Fixed Bugs With Both Legato and Sun Bug IDs

TABLE 1-1 contains descriptions of bugs or escalations (escal) with both Legato and Sun Bug IDs fixed in the 6.1 release of Solstice Backup software.

TABLE 1-1   Fixed in Solstice Backup Software, Release 6.1 *(1 of 3)*

Number	Description	Type
Legato Bug ID: LGTpa32496 Sun Bug ID: 4518960	A memory leak in the nsrd has been resolved.	Bug/escal
Legato Bug ID: LGTpa32258 Sun Bug ID: 4505623	Enabler code from previous version of Solstice Backup does not work correctly in release 6.0.	Bug/escal

TABLE 1-1   Fixed in Solstice Backup Software, Release 6.1 *(2 of 3)*

Number	Description	Type
Legato Bug ID: LGTpa30864 Sun Bug ID: 4462027	In a static/dynamic drive shared library environment, the jbverify command was incorrectly identifying shared drives in a jukebox and sometimes misidentified the drive order when used in DDS environments. The jbverify command has been modified to correctly identify shared drives in a DDS environment.	Bug/ escal
Legato Bug ID: LGTpa32112 Sun Bug ID: 4462769	Install does not assign right media type to DLT1 tape drives.	Bug/ escal
Legato Bug ID: LGTpa22483 Sun Bug ID: 4486220	Solstice Backup file permissions do not meet Sun standard.	Bug/ escal
Legato Bug ID: LGTpa30483 Sun Bug ID: 4462970	Solstice Backup release 6.1 will not install via jumpstart.	Bug/ escal
Legato Bug ID: LGTpa30618 Sun Bug ID: 4461393	The /etc/LGTOuscsi/inquire command did not display all tape devices attached to a system. This issue has been resolved. The Solstice Backup software now checks for HBAs with more than 16 targets.	Bug/ escal
Legato Bug ID: LGTpa30258 Sun Bug ID: 4496072	The nwadmin Pending window does not clear messages.	Bug/ escal
Legato Bug ID: LGTpa28408 Sun Bug ID: 4411164	When the Solstice Backup software reached the end of the tape label template number sequence, it did not wrap around to the start value. It continued to label with the last defined value producing an error.	Bug/ escal
Legato Bug ID: LGTpa30132 Sun Bug ID: 4459277	When immediate save and regular saves were interleaved, in some cases, the record number and the file number were incorrectly written into the media database record thereby corrupting the media database.	Bug/ escal
Legato Bug ID: LGTpa28108 Sun Bug ID: 4394226	mminfo could return incorrect query results when a -q savetime<=xx/xx/xx query constraint was specified.	Bug/ escal
Legato Bug ID: LGTpa28857 Sun Bug ID: 4483361	The volumes option of the recover command was not displaying all volumes required for recovery when the saveset spanned multiple volumes.	Bug/ escal
Legato Bug ID: LGTpa29112 (Duplicate of LGTpa28433) Sun Bug ID: 4431603	lus driver does not load on Solaris x86 version 7 or 8.	Bug/ escal

TABLE 1-1 Fixed in Solstice Backup Software, Release 6.1 *(3 of 3)*

Number	Description	Type
Legato Bug ID: LGTpa26056 Sun Bug ID: 4384936	nsrck does not delete orphan client index files.	Bug/ escal
Legato Bug ID: LGTpa30066 Sun Bug ID: 4458581	lus_add_fp_devs script fails if FC HBA buses are used and no tapes are present.	Bug/ escal
Legato Bug ID: LGTpa29613 (Duplicate of LGTpa29498) Sun Bug ID: 4461660, 4461606, 4461830, 4461314	nsrd to core dumps when running jbconfig on remote storage node.	Bug/ escal
Legato Bug ID: LGTpa27532 Sun Bug ID: 4454412	The recover program on NetWare reported no versions found message for valid entries.	Bug/ escal
Legato Bug ID: LGTpa30870 (Duplicate of LGTpa30383) Sun Bug ID: 4464249	Documentation error regarding directing Indexes to separate pools.	Bug/ escal
Legato BetaSphere Bug ID: BU-0062 Sun Bug ID: 4461742	How to Determine Device Mapping to Drives on L180/L700 Library. The fix for this bug is documented in the Solstice Backup 6.1 Release Notes, in Chapter 4, page 61, under the heading "How to Determine Device Mapping to Drives on L180/L700 Library."	Bug/ escal
Legato Bug ID: LGTpa31136 Sun Bug ID 4481873	Problem Installing NetWorker Modules with Solstice Backup because of incorrect package dependencies. The dependencies were corrected for Solstice Backup release 6.1.	Bug/ escal

Fixed Bugs With Legato Bug IDs

TABLE 1-2 contains descriptions of bugs or escalations (escals) fixed—and Requests for Enhancement (RFEs) implemented—in the 6.1 release of Solstice Backup software.

TABLE 1-2 Fixed in Solstice Backup Software, Release 6.1 *(1 of 4)*

Number	Description	Type
LGTpa12377	The pkgrm command on a Solstice Backup server failed to stop daemons. This problem has been resolved.	Bug/ escal
LGTpa16775	Certain license identifiers did not match. This problem has been resolved.	Bug/ escal

TABLE 1-2 Fixed in Solstice Backup Software, Release 6.1 *(2 of 4)*

Number	Description	Type
LGTpa16865	The `load_sleep` command was not functioning with an ADIC DAS silo. This problem has been resolved.	Bug/escal
LGTpa19085	Some recyclable volumes were being relabeled as blank. This problem has been resolved.	Bug/escal
LGTpa20845	When labeling a range of volumes from the GUI by changing the Starting With field from the default to another range, the first volume was labeled correctly, but subsequent volumes were labeled according to the Next field of the related label template resource. This problem has been resolved.	Bug/escal
LGTpa21685	In order to complete an NDMP directed recovery, certain client destination fields had to be changed from their normal Solstice Backup settings. This problem has been resolved.	Bug/escal
LGTpa21761	On HP-UX platforms, file timestamps were different when recovered by non-root users. This problem has been resolved.	Bug/escal
LGTpa21989	If the server client index in the `nwadmin` window indexes tab was selected to view the save sets of that client, it could take a long time to get the requested information back. While waiting, `nwadmin` and `nsrmmdbd` used up enough server CPU cycles to affect users who were trying to access a relational database. This problem has been resolved.	Bug/escal
LGTpa22680	Save set timeout was not respected for Legato NetWorker Module for Oracle backups. This problem has been resolved.	Bug/escal
LGTpa22869	The Sony AIT-2 drive was incorrectly being reported as a 4mm tape drive. This problem has been resolved.	Bug/esca
LGTpa24124	In HP-UX, the default device displayed during jukebox configuration displayed incorrectly. This problem has been resolved.	Bug/escal
LGTpa24533	The `nsrd` command would fail if the server hostname was upper case or mixed case. This problem has been resolved.	Bug/escal
LGTpa24923	References to "Emass Grau" were changed to "ADIC Emass Grau".	RFE
LGTpa25216, LGTpa25779	Save set recovery of NDMP backups that spanned more than one tape were hanging. This problem has been resolved.	Bug/escal
LGTpa25501	If no timeout value was configured in the *GROUP*.res file, the timeout value defaulted to the next midnight. This problem has been resolved.	Bug/escal
LGTpa25514	A group scheduled to do a skip-level backup was backing up "bootstrap" at level full. This problem has been resolved.	Bug/escal

TABLE 1-2 Fixed in Solstice Backup Software, Release 6.1 *(3 of 4)*

Number	Description	Type
LGTpa26209	The `nwbackup` GUI sometimes failed on Tru64 UNIX. This problem has been resolved.	Bug/escal
LGTpa26281	The `nsrjb -p` command option was failing if the slot range was greater than the number of drives. This problem has been resolved.	Bug/escal
LGTpa26595	There was an intermittent tape unload/eject issue with `nsrjb`. This problem has been resolved.	Bug/escal
LGTpa26688	A typographical error in a message generated by the `jbconfig` program was corrected.	Bug/escal
LGTpa26723	Tape devices were intermittently failing to write when the `nsrmmd` program on a storage node could not reconnect to a server after a particular RPC failed. This problem has been resolved.	Bug/escal
LGTpa26787	The `nsrjb` command could not access a configured device and was stopping during a backup. This problem has been resolved.	Bug/escal
LGTpa26847	Support for JNI PCI 64-bit fiber channel has been added.	RFE
LGTpa27124	Recovery was failing due to incorrect conversion of a tape offset in the media database. This problem has been resolved.	Bug/escal
LGTpa27405	File indexes now require less space for the same number of records.	RFE
LGTpa27510	The online Help for the Backup Command field and `savepnpc` needed updating. It has been updated.	Bug/escal
LGTpa27520	During a command-line recovery, it was possible to change directories (cd) to an invalid directory. This problem has been resolved.	Bug/escal
LGTpa27895	File recall was failing when a user was a member of more than 16 groups. This problem has been resolved.	Bug/escal
LGTpa28000	Temporary 2fh files were not being removed from `/nsr/tmp` after a backup completed. This problem has been resolved.	Bug/escal
LGTpa28055	The `jbconfig` and `inquire` commands were failing with a "bailing out" error message. This problem has been resolved.	Bug/escal
LGTpa28168	Non-root users were not allowed to have `chown` privileges on HP-UX if a recovery operation was to be performed. This problem has been resolved.	Bug/escal
LGTpa28814	Modifying expired browse and retention policies on converted Solstice Backup 5.x save sets could result in a core dump. This problem has been resolved.	Bug/escal
LGTpa28877	Environment variables (for tape devices) containing hyphens were not passed correctly in the Bourne (sh) or Korn (ksh) shells. This problem has been resolved.	Bug/escal

TABLE 1-2   Fixed in Solstice Backup Software, Release 6.1 *(4 of 4)*

Number	Description	Type
LGTpa29186	Support for Sun's new fiber-channel drivers (including support for FC-TAPE) was added.	RFE
LGTpa29674	NDMP directed recovery from a local client to a third- party client failed. This problem has been resolved.	Bug/ escal
LGTpa29793	The file ndmpjbconf has been added to the list of files to be put into lgtodrvr on Linux86 to support the NetWorker NDMP module.	RFE
LGTpa29835	Solstice Backup software can now use SmartMedia resources on AIX.	RFE
LGTpa30026	The handle count and virtual memory for nsrexecd were growing at a rapid rate, indicating that they were not being freed properly. This problem has been resolved.	Bug/ escal

Installation and Upgrade Notes

This section includes topics relevant to installing and upgrading the Solstice Backup software.

Configuration Script for Fiber Channel Devices

If you have fiber channel devices in your Solstice Backup configuration, you must reconfigure the fiber channel jukebox and fiber channel tape devices after applying patch 111962-xx. To reconfigure the fiber channel devices, run the following command after you have installed the patch:

```
/etc/LGTOuscsi/lus_add_fp_devs
```

Changes Resulting from the Improved Save Stream File Format

An improved save stream file format was implemented in Solstice Backup release 6.0. This format makes possible many enhancements and features, such as NDMP, large save sets, and the improved media database.

Because this format is significantly different, some backward incompatibility results if Solstice Backup 6.0 and later clients or servers are combined with pre-6.0 Solstice Backup clients or servers. Limitations encountered as a result of the new file format can be resolved by upgrading to the Solstice Backup 6.0 or later client and server software.

The following sections describe areas in which a file format mismatch would be evident.

Continued Save Sets Eliminated

The 64-bit format eliminates the need for continued save sets when saving data. To use the larger save sets, the clients and servers must both be running Solstice Backup release 6.0 or later software. As a result of this change, the Save Set Cutoff value in the Server resource is no longer used unless the client software is earlier than Solstice Backup release 6.0. A value in this field is ignored if you are backing up a Solstice Backup 6.0 or later client.

Save Set Names Longer Than 64 Bytes

The 64-byte size limitation on save set name lengths has been removed. You can create save set names up to 255 bytes long. To use the longer names, the clients and servers must both be running Solstice Backup release 6.0 or later software.

Saving and Recovering between Different Releases

Using the Solstice Backup release 6.0 save stream format, you can easily recover data saved using either the release 5 or release 6 format. The format change is transparent to you unless you have data saved in the later format and try to recover it using the earlier format. The earlier format is unable to recover the more recent format. TABLE 1-3 lists the possible combinations and their results.

TABLE 1-3   Save Stream Matches and Mismatches

Saved from and to:	Recovered by:	Recovery Successful?
6.x client to 6.x server or storage node	6.x client	Yes
6.x client to 5.x server or storage node[1]	6.x client	Yes
5.x client to 6.x server or storage node	6.x client	Yes
5.x client to 5.x server or storage node	6.x client	Yes
6.x client to 6.x server or storage node	5.x client	No
5.x client to 6.x server or storage node	5.x client	Yes
6.x client to 5.x server or storage node[1]	5.x client	Yes
5.x client to 5.x server or storage node	5.x client	Yes

1. Solstice Backup does not support configurations in which the backup server or storage node is at a lower version than any of its clients, regardless of platform.

You might receive the following error message under either of the following conditions:

```
recover failed:Error receiving files from NSR server '<server>'
```

- Attempting to recover files backed up by Solstice Backup release 6.0 or later software using Solstice Backup release 5.x. The 5.x client does not recognize the 6.x format.
- The parallelism value is set too high, causing the recovery session to fail.

To solve this problem, do one of the following:

- Upgrade the client from Solstice Backup release 5 software to release 6.0 or later software.
- Recover the file to a client that has Solstice Backup release 6.0 or later software installed.
- Decrease the parallelism value setting in the Client or Server resource. This should improve recovery performance.

Backward Compatibility of Storage Volumes

Once a Solstice Backup release 6.0 or later server or storage node is used to write a tape volume, that volume can no longer be read by the server or storage node from an earlier release. This includes labeling operations, as well as writing save sets. There is no problem when pre-6.0 volumes are read by pre-6.0 storage nodes, even when the server controlling the node is running Solstice Backup release 6.0 or later (a common situation if, for example, a storage node has its own jukebox).

The restriction also applies to a shared jukebox configuration, where tape volumes might readily be read by any host (server or storage node) connected to the jukebox. To avoid the problem, once the server—or any storage node connected to a 6.0 or later server—in the configuration is upgraded to Solstice Backup 6.0 or later software, upgrade all components to at least Solstice Backup 6.0 software.

All servers and storage nodes running Solstice Backup 6.0 or later software can read from, and append to, tape volumes from all releases of the Solstice Backup software.

Upgrading Storage Nodes

Storage nodes running release 6.0 or later do *not* work with servers running releases earlier than 6.0. Upgrade servers before upgrading storage nodes.

Administering Solstice Backup 6.x Servers from Earlier Clients

You can create Solstice Backup resources (such as directives and archive annotations) with Resource Allocation Protocol (RAP) values that are larger than 1 KB. The 1 KB size limitation has been removed. Clients running the Solstice Backup 5.5.x administration program, however, do not support resource values larger than 1 KB. Consequently, to use the larger resource values, clients and servers must both be running Solstice Backup release 6.0 or later software. Once the server contains resource values larger than 1 KB, it is important that all administration programs be upgraded.

Backup and recovery operations are not affected by this incompatibility. Only administration programs are affected.

New or Updated Licenses

The following section describes new or updated licences in Solstice Backup release 6.1.

Update Enablers

If you have an enabler for an earlier release of the Solstice Backup software and are installing this release as an update, you received the update enabler code on the Enabler Certificate. The update enabler brings forward all the licenses for which you have already authorized. Enter the update enabler code as you would any other enabler code and do *not* remove any of the previously authorized codes.

Note – Customers who have Sun support contracts for Solstice Backup software automatically receive update enabler codes at no additional charge.

New Cluster Client Connection Licenses

With Solstice Backup 6.0 and later, a Cluster Client Connection license is required in order to back up a client (physical host) belonging to a cluster. Releases of Solstice Backup prior to 6.0 did *not* require this license and allowed the backup of individual nodes of a cluster using the Client Connection licenses. In 6.0 or later, backups of clustered clients will fail without the Cluster Client Connection license.

Note – Solstice Backup 6.1 is supported in the Sun Cluster 3.0 environment but is not supported in the Sun Cluster 2.2 environment. However, Solstice Backup 6.0 software is supported in the Sun Cluster 2.2 environment.

If you are updating from a version of Solstice Backup prior to 6.0 and wish to back up the hosts of any cluster framework, you must purchase the appropriate number of Cluster Client Connection licenses. The licenses can be used with either Network Edition or Power Edition, but not with Workgroup Edition.

Two Cluster Client Connection licenses are standard in the Power Edition of Solstice Backup.

The new Cluster Client Connection licenses are:

- LSBIS-610-2114 Cluster Client Connection License for UNIX, Quantity 1
- LSBVS-610-2115 Cluster Client Connection License for UNIX, Quantity 5
- LSBYN-610-2110 Cluster Client Connection License for NT, Quantity 1
- LSBYN-610-2111 Cluster Client Connection License for NT, Quantity 5
- LSBIS-610-7058 Cluster Client Connection License for Linux, Quantity 1
- LSBIN-610-7059 Cluster Client Connection License for Linux, Quantity 5

About Servers with Multiple Domains

On servers that have multiple domains, each domain is considered to be a separate system. Since each Solstice Backup license applies to a single system, each domain would need its own Edition license to be able to act as a backup server.

A Backup Server or Storage Node addresses the needs of one domain in a multiple-domain host system. Libraries can be shared across domains. Drives can be shared across domains with the optional Dynamic Drive Sharing license.

The following table shows the kinds of licenses required for a domain based on the role that is configured for that domain.

TABLE 1-4 Licensing Servers with Multiple Domains

Domain's Role	License Needed
Backup Server	NetWork, Power, Storage Node or SAN Storage node Server License
Storage Node	Storage Node or SAN Storage Node License[1]
Backup Client	Client Connection License
Backup Server or Storage Node(s) sharing library across tape drives	NetWork or Power Edition, Storage Node(s) or SAN Storage Node(s)[1] (using inherent Library Sharing feature)
Backup Server and Storage Node(s) dynamically sharing libraries and all tape drives	NetWork or Power Edition, Storage Node(s) or SAN Storage Node(s);[1] (with a Dynamic Drive Sharing license for each tape drive)

TABLE 1-4   Licensing Servers with Multiple Domains

Domain's Role	License Needed
Multiple Backup Servers sharing drives across domains in multiple datazones	Network or Power Edition for each Backup Server; GEMS Smart Media[2]
Multiple Backup Servers managed through one Backup console	GEMS Console[2]

1. The license must match that of the backup server edition.

2. GEMS Smart Media and GEMS Console work with Solstice Backup 6.1 but are currently available only from Legato systems. Contact sun@legato.com for pricing and ordering information.

Licensing Solstice Backup for Linux

As of Solstice Backup release 6.0, Solstice Backup for Linux client software is licensed as a separate platform.

Any new or existing customers who wish to have Linux client support will need to purchase the ClientPak for their Linux product; the ClientPak for UNIX will no longer include Linux support. To purchase the ClientPak for Linux or request more information on the UNIX ClientPak, contact your Sun sales representative or service provider.

To continue receiving support on Linux clients while you obtain a permanent ClientPak for Linux enabler, use the following 45-day temporary ClientPak for Linux enablers:

- f67378-9f0226-6818bc
- d25754-b32e0a-794f98

New Enterprise Application Module for Oracle 3.0

Solstice Backup releases after version 5.5 provide a choice of two modules for Oracle support rather than one. If you are updating from 5.5 to Solstice Backup 6.0 or later or purchasing Solstice Backup 6.0 or Solstice Backup 6.1 for the first time, you must chose the appropriate Enterprise Application Module (EAM) for Oracle support. Your decision should be based on the version of Oracle to be protected. See the following table for reference on choosing an appropriate EAM "Enterprise Application Module Options for Oracle 3.0" :

TABLE 1-5 Enterprise Application Module Options for Oracle 3.0

Versions of Solstice Backup	Oracle 7.x	Oracle 8. 8i
5.5.1 without patch	LSBIS-600-2151 or LSBIS-610-2151	LSBIS-600-2151 or LSBIS-610-2151
5.5.1 with patch 108884-02 or later	LSBIS-600-2151 or LSBIS-610-2151	LSBIS-600-2154 or LSBIS-610-2154
6.0	LSBIS-600-2151 or LSBIS-610-2151	LSBIS-600-2154 or LSBIS-610-2154
6.1	LSBIS-600-2151 or LSBIS-610-2151	LSBIS-600-2154 or LSBIS-610-2154
Note: LSBIS-600-2151, LSBIS-610-2151, LSBIS-600-2154 and LSBIS-610-2154 are backward-compatible to earlier versions of the Solstice Backup software.		

The new Oracle 3.0 module was released by Legato Systems between the Solstice Backup 5.5 and Solstice Backup 6.0 releases. The EAM for Oracle 3.0 was introduced in Solstice Backup 6.0 (announced in October 2000). For more information, consult the Legato Software Compatibility Guide at:

```
http://www.legato.com/resources/compatibility
```

Note – The update enablers (for example, Update/7) do *not* apply to the Enterprise Application Modules (or Database Modules). An update enabler applies when upgrading the Solstice Backup server software from an earlier version.

New Licenses for SAN Storage Nodes

The licenses for SAN Storage Nodes are available for purchase in quantities of five (5). A license for a quantity of five comes with five individual enabler codes. To enable the SAN Storage Nodes functionality, you can distribute these enablers to Solstice Backup servers in different data zones.

The new licenses are:

■ SAN Storage Node for UNIX, Network Edition, Quantity=5

■ SAN Storage Node for UNIX, Power Edition, Quantity=5

■ SAN Storage Node for Windows 2000/NT, Network Edition, Quantity=5

■ SAN Storage Node for Windows 2000/NT, Power Edition, Quantity=5

■ SAN Storage Node for Linux, Network Edition, Quantity=5

■ SAN Storage Node for Linux, Power Edition, Quantity=5

On the enabler certificate there are five enabler codes. To enable and register this option, you must do the following:

■ Enable and register your product; be sure to enter each enabler code.

■ When you receive the permanent authorization codes for each enabler code that you registered, enter the permanent authorization codes. For example, if you have enabled and obtained authorization keys for five servers, enter the authorization key for each one.

Note – For detailed instructions on how to enable and register your product and to obtain permanent authorization keys, refer to the *Solstice Backup 6.1 Installation Guide.*

Important Notes and Tips

This section describes important notes and helpful tips, including configuration tips.

Apply Latest Compaq Patches If System Has Large File Systems Mounted

If you have Solstice Backup release 6.0 or later installed on Compaq Tru64 UNIX release 5.x, scheduled save sets might fail if you have large file systems (of or about 1 terabyte in size) mounted.

To complete scheduled save sets, ensure that the Tru64 version 5.1 Patch Kit 3, or later, is applied. For further information regarding this patch, contact Compaq.

Backing Up Large Numbers of Files on HP-UX Requires Increasing maxdsiz Parameter

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. Do this by changing the maxdsiz parameter to 2 GB using the `sam` utility.

To change the maxdsiz parameter:

1. As `root`, 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 in an NDMP Environment with Data ONTAP 5.3.x

For information about this topic, see .

Block Size Differences on Storage Node Can Result in Failure of Recovery

SCSI hardware limitations differ between UNIX and Microsoft Windows platforms. This can lead to compatibility problems. For example, if you have a device defined in UNIX that is physically attached to a Microsoft Windows host bus adapter (HBA), it is possible to define a block size greater than that allowed by the hardware on which Microsoft Windows is running. This can lead to I/O errors in both write and read states on the device. If you use both platforms, you must determine a common block size acceptable to both platforms.

Bus Reset Can Rewind Tape on Devices Shared by Multiple Computers

Tape devices shared by more than one computer can, at unpredictable times, experience bus resets from any of the computers. These reset commands can cause a tape on a shared bus (such as SCSI or fiber channel) to rewind. The results can include:

- Tapes that are prematurely treated as full
- Corrupted data on tapes

System configurations that do *not* properly protect tape devices shared by more than one computer can experience these bus resets. Some switching hardware can be configured to protect tape devices from resets; certain operating systems include built-in protection (that can be turned on by the user) against stray bus resets. To determine whether your switch or operating system includes such protection, and to learn how to use it, refer to the documentation provided by the manufacturer or contact the manufacturer.

Therefore, configuring a tape device in a shared SCSI or fiber channel environment without using either a hardware switch or appropriate operating system solution that handles the bus reset issue is not supported by Solstice Backup software.

Apply Proper Point Patch to Correct jbverify Defect (Legato Bug ID: LGTpa31520, Sun Bug ID: 4523123)

The Solstice Backup software contains a defect that causes jbverify to fail while checking the second device in a jukebox.

If you are running Solstice Backup 6.1 with patch 111962-xx, apply point patch 111971-xx to correct this defect.

Note – Solstice Backup 6.1 with patch 111962-xx is a prerequisite for point patch 111971-xx.

If you are installing the Solstice Backup 6.1 L10N release, you must apply point patch 111970-xx to correct this defect.

Known Limitations

This section describes known limitations in Solstice Backup. When appropriate, both Legato and Sun Bug IDs are included.

Limitations in Feature Support

Feature-related limitations for Solstice Backup release 6.1 with Patch 111962-xx are listed below:

Recovering Several Subdirectories May Cause Solstice Backup to Hang in an NDMP Environment (Legato Bug ID: LGTpa32414)

When recovering several NDMP subdirectories using the nwrecover command, the Solstice Backup software might hang if you do not shutdown and restart nwrecover after the first successful subdirectory has been recovered.

After the first subdirectory has been recovered, exit and start nwrecover again and then recover the remaining directories.

Unable to Recover Files in a Subdirectory from Multiple Tapes using nwrecover in a NDMP Environment (Legato Bug ID: LGTpa32433)

When recovering files and subdirectories from multiple tapes, the nwrecover command only recovers files from the first tape volume successfully in a NDMP environment.

Recovery of subdirectories and files on subsequent tapes can be recovered one file at a time. You might have to reboot if the error occurs.

Note – This problem only occurs with EMC Celerra version 2.2.xx.x.

Element Status Feature Must be Manually Enabled for Jukeboxes That Support the Feature (Legato Bug ID: LGTpa26003, Sun Bug ID: 4306035)

When a jukebox supports the element status feature but that feature is not enabled in the Autochanger resource of the Solstice Backup Administrator program, nsrjb -I and nsrjb -H will *not* work properly.

When using jukeboxes that support the element status feature, make sure that this feature is enabled in the Autochanger resource of the Solstice Backup Administrator program before using nsrjb.

Using Solstice Firewall v4.1sp3 and Non-Sync Packets (Legato Bug ID: LGTpa30256, Sun Bug ID: 4427943)

Because Solstice Backup uses non-sync packets to perform backups, when using the Solstice Firewall v4.1sp3, these packets will be dropped unless you specify that the rule base should allow them.

To specify that the rule base should allow non-sync packets, uncomment the following line in the $FWDIR/lib/fwui_head.def file (where $FWDIR is the installation directory of your firewall software; the default location is /etc/fw):

```
define ALLOW_NON_SYN_RULEBASE_MATCH
```

jbexercise -D and -d Options Do Not Work (Legato Bug ID: LGTpa30555, Sun Bug ID: 4462022)

The `-D` and `-d` options for the `jbexercise` program do *not* work. The workaround is to use `jbexercise` with the `-m` *other* option rather than specifying the jukebox model.

Incremental Backups after Renaming Files and Directories (Legato Bug ID: LGTpa00800)

When recovering a directory that has been renamed since the last full backup (and only incremental backups have been run since then), you might find that files that you expect to find under the renamed directory are not there.

To recover those files, either perform a full backup before starting the recovery, or do the following:

1. Shut down the Solstice Backup daemons on the server and client.

2. Set the environment variable `NSR_RENAME` to `Yes` before starting the client.

3. Set the environment variable `NSR_NAMEONLY` to `No` before starting the server.

4. Start the recover session.

Re-Opening a Closed Save Set Window during a Recovery (Legato Bug ID: LGTpa27875)

Leave the Save Set Recover and the Save Set Recover Status windows open while a `recover` operation is in progress. If either of these windows is closed before the recovery completes, the `nwadmin` program causes a core dump if the window is then re-opened.

The nsrjb Command Tries to Verify Tape Label with No Tape in Drive (Legato Bug ID: LGTpa27603)

When the `nsrjb` command requests that a tape be placed into a drive in an IBM 3594 two-drive silo, the silo attempts to move the tape to the drive, then reports that the drive has failed and returns the tape to its slot.

At this point, Solstice Backup software attempts to verify the tape's label, but since there is no tape in the drive, the software reports an input/output error and attempts to remove the tape from the drive. An error message similar to the following then appears:

```
Jukebox 'IBM_3494' failed cannot find slot '-9999 in operation
slots list.
```

The following workaround is available:

1. Disable the drive that failed.

   In the `nwadmin` GUI:

   a. From the Media menu, select Devices.

   b. Highlight the device that you want to disable.

   c. Select No next to the Enabled attribute.

2. Shut down the Solstice Backup software.

3. Make a backup copy of the `nsrjb.res` file. Then edit the original file to remove the offending -9999 entry.

4. Restart the Solstice Backup software.

Problem with Instances Listbox in GUI (Legato Bug ID: LGTpa25616)

If you select an instance in the Instances Window of the Save Set Recover dialog box and double-click it, or if you select multiple instances and click the Details button, the information that appears in the Save Set Recover Details window might not pertain to the instance you had expected. To avoid this problem, do not double-click instances, and only click the Details button after a single instance has been selected.

Mounting a Volume during a Heavy Server Load (Legato Bug ID: LGTpa25762) (Resolved in Solstice Backup 6.1 with Patch 111962-xx)

When a tape mount is attempted while the server load is heavy, you might see the following message:

```
volume_name volume id volid not found in media index
```

If, however, you manually query the media index by entering, for example:

```
mminfo -vm volume_name
```

The command output shows that the volume does exist with the stated volume ID.

Subsequent tape mount requests by Solstice Backup software to service a volume might indicate that the volume ID is not found in the media index. This problem continues until the server load is reduced.

To reduce the server load, reduce the parallelism.

Adequate Licensing Required to Administer Legato GEMS (Legato Bug ID: LGTpa27548) (Resolved in Solstice Backup 6.1 with Patch 111962-xx)

When using Legato GEMS, ensure that you have enough licenses for the nodes you intend to administer. If you do not have enough licenses, the Sun license management daemon (lgtolmd) might fail.

Mismatching Host Prefixes in Jukeboxes (Legato Bug ID: LGTpa29474)

If, during jukebox configuration, two slightly different forms of the hostname are entered for different jukeboxes on the same host, it is possible that you could lose the ability to access one of the jukeboxes. It is not likely that this would happen, but because the `nsrd` program does not currently check jukebox names during device creation, the problem can occur in the following situation:

- A storage node has a jukebox (jukebox_1) with no storage devices defined on that particular storage node (assume that the tape devices in the jukebox were defined on a different storage node on the same server).

- When jukebox_1 was configured using the `jbconfig` program, the default form of the local hostname (for example, `mercury.companyname.com`) was not used; a shorter name (`mercury`) was selected instead.

- A second jukebox (jukebox_2) that *does* contain local tape devices is now configured. This time, the default form of the hostname (`mercury.companyname.com`) is used.

- The local hostnames for the two jukeboxes now do not match (`mercury` and `mercury.companyname.com`).

When such a hostname mismatch occurs, `jukebox_1` (the one with no local storage devices) is no longer accessible.

Recommended Maximum Client/Server Parallelism Values (Legato Bug ID: LGTpa28774)

To ensure that the Solstice Backup software backups complete successfully, do not set the client parallelism attribute higher than the server parallelism. The current maximum value for server parallelism is 512.

Only 64 Devices Supported in a Single Data Zone (Legato Bug ID: LGTpa29556)

Although the jbconfig program in Solstice Backup release 6.1 supports a maximum of 256 devices, the underlying operating system imposes its own, additional limits on the number of supported devices and processes (based on available hardware resources). Refer to your operating system's documentation for more information regarding such limits.

Setting Reserve/Release Bit for Correct Operation of DDS (Legato Bug ID: LGTpa30565)

With Dynamic Drive Sharing (DDS) enabled on your computer, you might encounter a situation in which a shared drive does *not* mount or unmount a tape (or other media) correctly. This error might be caused by the reserve/release bit setting in the operating system.

To correct this situation and enable the shared drive to mount/unmount media correctly, you must reset the reserve/release bit setting for the operating system. See the following sections for specific information about how to reset the reserve/release bit setting for certain operating systems.

Solaris

On a Solaris operating system:

1. Access the `st.conf` file.

2. To enable DDS successfully, add the value of `0x2` to the fourth field of the data property line for the device that is to be shared. For example:

```
ST_NO_RESERVE 0x20000
```

3. After saving the change to the file, enter the following commands:

```
# drvconfig -i st
# tapes
```

For more information, refer to the Tape Configuration section of the `st` man page.

AIX

On an AIX operating system:

1. Through the `smit` or `smitty` interface, select Devices>Tapes.

2. Change the value for the attribute RESERVE/RELEASE support from no to **yes**.

Testing Tapes Using the tapeexercise or tapeexer Command (Legato Bug ID: LGTpa29268)

The following points apply to testing a tape using the `tapeexercise` (Microsoft Windows) or `tapeexer` (UNIX) command:

■ Do *not* use the `tapeexercise` or `tapeexer` command on any tape containing data that you will need later, because the tests overwrite data currently on the tape.

■ After testing is completed, re-label any tape that you have used in testing. If the tape is not re-labeled, it cannot be used for subsequent backups or recoveries.

■ Using barcode-labeled tapes for testing is not recommended. It is more difficult to re-label a tape that has a barcode label, because the barcode label is linked to the volume ID, which is in the media database.

Error Messages During Automatic Space Recover Can Be Ignored (Legato Bug ID: LGTpa30741)

If an automatically scheduled recover space operation (on a file-type device configured for staging) is unable to take place, an error message appears, describing why the operation failed. For example:

```
06/20/01 12:06:02 nsrd: recover space operation on device
[device-name] failed, device '[device-name]' is busy.
```

Note – This message can safely be ignored, since the automatic space recover operation will run at the next specified interval, assuming that the device is available at that time.

Using Online Help in Conjunction with the Details View (Legato Bug ID: LGTpa31010)

The `nwadmin` GUI currently dumps core if you choose the Details view from a resource window—such as Devices—while a Help window is open, but no Help topic is selected (highlighted).

Since the default is for a Help topic to be selected, the only way to encounter this bug is specifically to deselect the highlighted Help topic and then choose the Details view in the resource window.

The workaround is to avoid choosing the Details view if the Help window is open *and* the highlighted Help topic has been deselected. A deselected topic shows no highlighting, but is simply surrounded by a line of dash marks (- - -).

Limitations in Platform Support

Platform-related limitations are listed below.

Using the `jbconfig` Command with Sun StorEdge RAID Manager (Legato Bug IDs: LGTpa24044, 27298, Sun Bug ID: 4358034)

When Sun StorEdge™ disk arrays are used in conjunction with the Sun StorEdge RAID Manager, the `inquire` program might take an unexpectedly long time to complete. The `jbconfig` command uses the `inquire` program to obtain information about autochangers that are attached to the system.

A possible workaround for this problem is to stop the RAID Manager daemons before running `jbconfig`. One disadvantage of this workaround is that RAID Manager features, such as disk failover protection, are not operational during this period. However, `jbconfig` is usually run only during installation and initial configuration of the Solstice Backup software, and subsequently when adding autochanger hardware to the system. Often these tasks are performed during scheduled system maintenance periods when shutdown of the RAID Manager software is acceptable. If you want the `jbconfig` program to complete rapidly, then it is best to use `jbconfig` during scheduled system maintenance periods after stopping the RAID Manager daemons.

Another possible workaround for this problem is to allow the `jbconfig` command to execute with the RAID Manager daemons still running. If this method is used, `jbconfig` may take a long time to complete, and this may cause unacceptable system performance issues.

Another workaround is to perform the following procedure:

1. In /etc/driver_classes, delete the following line:

```
rdnexus scsi
```

This will prevent the framework from enumeration scsi target nodes under rdnexus.

2. In rdriver.conf, change:

```
class="scsi"
```

to

```
parent="rdnexus"
```

This is will cause rdriver nodes to be created under rdnexus, but not under physical hba drivers.

This change will only last until the next boot -r or hot_add.

Optical Disk Restriction on AIX (Legato Bug ID: LGTpa21747)

If you run Solstice Backup software under AIX 4.3.2 and attempt to perform a backup using either an IBM® 3995 or HP-C1160 optical library that spans more than one volume, the backup operation fails after the first volume is filled.

UFS Files Having Inode 4 or 5 Recovered to an AdvFS Filesystem on Tru64 UNIX (Legato Bug ID: LGTpa25810)

Do *not* attempt to recover a UFS file that has an inode of either 4 or 5 to an AdvFS filesystem on Tru64 UNIX. Such an attempt fails, because AdvFS user quota files are already assigned inode 4 and group quota files are already assigned inode 5. These AdvFS quota files cannot be overwritten. (Solstice Backup software uses a special algorithm to save and recover AdvFS quota files.)

To recover a UFS file having either inode 4 or 5 to an AdvFS filesystem, do the following:

1. Recover the file to a UFS filesystem.

2. Enter the `cp` command to copy the file to the desired location in the AdvFS filesystem.

Compaq Patches Required for Solstice Backup Release 6.0 or Later on Compaq Tru64 UNIX 5.x (Legato Bug ID: LGTpa27333)

If you run Solstice Backup release 6.0 or later on Compaq Tru64 UNIX release 5.1 or 5.0a, you must install the appropriate Compaq patch to ensure that save sets scheduled by Solstice Backup software complete successfully:

For Solstice Backup release 6.0 or later on:

■ Compaq Tru64 UNIX 5.1, install Patch 399.
■ Compaq Tru64 UNIX 5.0a, install Patch 286.

For more information regarding either of these two patches, contact Compaq.

Support for AIX5L Version 5.1 32-bit Kernel Mode Only (Legato Bug ID: LGTpa31332)

Solstice Backup software supports AIX5L Version 5.1 in 32-bit kernel mode. Solstice Backup software does *not*, however, support AIX5L Version 5.1 in 64-bit kernel mode.

Support for Bull and AIX SAN-Certified Solution (Legato Bug ID: 28460) (Resolved in Solstice Backup 6.1 Patch 111962-xx)

The fix is available on the ftp site but is *not* required when running Solstice Backup 6.1 with patch 111962-xx:

- Support for a Bull Fiber Channel (FC) environment when running Solstice Backup release 6.1 on AIX 4.3.3 is now available.

- A special Bull FC driver must be applied to the AIX 4.3.3 operating system. Special AIX file sets that support Fiber Channel must also be updated for an AIX 4.3.3 system.

- A new `cfglus` library has been created that recognizes fiber-connected devices on AIX when using the Bull driver *only*.

- Bull's FC implementation created new device objects /dev/fcp*.

- The `lus` configuration method has been changed to recognize the /dev/fcp* objects.

Requirements

- Solstice Backup release 6.1

- AIX 4.3.3

- AIX Base Operating System filesets updated to support Fiber Channel (FC)

- The updated `cfglus` library

- Bull Fiber-Channel driver installed, with the following firmware version minimums:

  - Required base drivers:

    devices.pci.df10e51a.diag   1.3.3.6

    devices.pci.df10e51a.rte   1.3.3.23

    devices.pci.df10e51a.ucode 1.3.3.0

  - With LP7000 HBA, add:

    devices.pci.df1001f7.diag   1.3.3.0

    devices.pci.df1001f7.rte   1.3.3.3

    devices.pci.df1001f7.ucode 1.3.3.6

  - With LP8000 HBA, add:

    devices.pci.df1001f8.diag   1.3.3.0

    devices.pci.df1001f8.rte   1.3.3.2

    devices.pci.df1001f8.ucode 1.3.3.6

Note – At the time of publication, patches (PTFs) for Fiber Channel drivers could be downloaded from: `http://www-opensup.bull.com`.

To download patches for Fiber Channel drivers:

1. Contact Bull directly for this driver and for additional information about it.

Caution – This fix should only be applied to Solstice Backup release 6.1. It should *not* be used on Solstice Backup release 6.0.x or 6.1 with patch 111962-xx.

2. Download the following `gzip` file from:

```
ftp://ftp.legato.com/pub/NetWorker/Updates/AIX/
LGTpa28460.tar.gz
```

3. Review the `LGTpa28460.readme` included at the above location for required information for installing and troubleshooting this implementation.

4. Review the SAN Solutions Compatibility information at the following location if you have questions regarding compatibility:

`http://www.legato.com/support/documentation/compatibility/`

Problem Creating Start/Stop Scripts When Installing SuSE 7.1 (Legato Bug ID: LGTpa30719)

To install the Solstice Backup software on the computer being designated as the Solstice Backup client:

1. For SuSE 7.1 only, create the `rc.d/init.d` directory in the `/etc` directory.

   For example:

```
# mkdir rc.d
# cd rc.d
# mkdir init.d
```

2. Install the Solstice Backup software:

    a. Enter the NWinstall command at the command prompt:

```
# ./NWinstall
```

    b. At the prompt, enter the corresponding number for the option to install the lgtoclnt package. Installing this package temporarily requires 70 MB of free space on the client computer.

3. (Optional) At the prompt, enter the appropriate number for the option to install the man pages (lgtoman). When you have added all the applicable packages and the prompt is displayed, exit the installation script.

4. Copy the Solstice Backup start/stop script (networker start) from the etc/rc.d/init.d directory to etc/init.d.

5. Remove the following two directories:

   – etc/rc.d/init.d

   – etc/rc.d

6. Create a symbolic link from:

   /etc/init.d/networker

   to:

   /etc/init.d/rc0.d/K05networker, /etc/init.d/rc3.d/S95networker, /etc/init.d/rc5.d/S95networker.

   For example:

```
# ln -s /etc/init.d/networker /etc/init.d/rc0.d/K05networker
# ln -s /etc/init.d/networker /etc/init.d/rc3.d/S95networker
# ln -s /etc/init.d/networker /etc/init.d/rc5.d/S95networker
```

7. After installation, you can change the list of trusted Solstice Backup servers that have access to the client's data by editing the /nsr/res/servers file. For detailed instructions on how to change the Solstice Backup servers that are authorized to access the client, see the section titled "Changing the Solstice Backup Servers with Access to a Client" in the *Legato NetWorker Installation Guide, Release 6.1, Linux Version.*

If the output does not list `nsrd` and `nsrexecd` as current processes, run the following command to start the Solstice Backup daemons on the Solstice Backup server:

```
# /etc/init.d/networker start
```

Cannot Perform Relocated Install on Red Hat 7.1 (Legato Bug ID: LGTpa31230)

For Red Hat Linux release 7.1 only, you must upgrade the following files available with the Red Hat Installation Media (CD 2) or at `http://www.rpmfind.net`:

- `rpm-build-4.0.3-0.66.i386.rpm` package
- `rpm-devel-4.0.3-0.66.i386.rpm` package
- `rpm-4.0.3-0.66.i386.rpm` package
- `glibc-2.2.3-14.i386.rpm` package
- `glibc-common-2.2.3-14.i386.rpm` package
- `glibc-devel-2.2.3-14.i386.rpm` package

Installation Fails on Caldera 3.1 Due to Missing Dependency (Legato Bug ID: LGTpa31607)

To install the Solstice Backup server software properly:

1. Save the original files described in the *Legato NetWorker Installation Guide, Release 6.1, Linux Version*, prior to the server installation step if you want to keep a copy of the current configuration.

2. Create the following symbolic link:

```
# ln -s /usr/lib/libtermcap.so libtermcap.so.2
```

3. Enter the following command:

```
# rpm -ivh --nodeps lgtoclnt-6.1-1.i386.rpm
```

4. Enter the NWinstall command:

```
# ./NWinstall
```

Continue the installation procedures as described in the *Legato NetWorker Installation Guide, Linux Version, Release 6.1*.

Save Group Failure on Standalone Devices on IRIX (Legato Bug ID: LGTpa31773)

On an IRIX platform, save groups initiated from a Solstice Backup server where a standalone device is the only device available and the automedia management feature is turned on will fail once the device is no longer operational (is disabled after the maximum number of consecutive errors has been exceeded).

If this happens, turn the automedia management feature off, and label the tape manually.

Documentation Errata

This section contains corrections or additions to the Solstice Backup release 6.1 documentation.

Installing to a Default Location (Solstice Backup Install Guide, page 23)

Ignore the following Note:

Note – To install and run any version of the non-english localized Solstice Backup software, release 6.0.1L—you must first install the localized English language release of Solstice Backup 6.0.1L. Once this release of the software is installed, you can install the required localized version of Solstice Backup. If you are currently running the localized English 6.0.1L Solstice Backup software, you can install the required localized software.

Also, ignore the following sentence under the Note:

Solstice Backup release 6.0.1L is the most recent localized version.

Update Enablers for Solstice Backup 6.1 (Solstice Backup Install Guide, page 36)

Replace *Table 3-1: Update Enablers for Solstice Backup 6.1* on page 36 with the following updated table:

Solstice Software Release	Update Enabler Required
Release 6.0 to 6.1	no update enabler required
Release 5.5/5.5.1 to 6.1	Update/7
Release 5.1/5.1.1 to 6.1	Update/7
Release 4.x to 6.1	not supported

Updating a Solaris Master Server from Solstice Backup Release 5.5.1 to 6.1 (Solstice Backup Install Guide, page 40-41)

After Step 8, the following note should be added:

Note – The remaining steps apply only to customers with maintenance contracts. Non-maintenance contract customers should proceed by using the enablers you have ordered.

Updating a Solaris Master Server from Solstice Backup Release 5.1/5.1.1 to 6.1 (Solstice Backup Install Guide, page 43-44)

After Step 8, add the following note:

Note – The remaining steps apply only to customers with maintenance contracts. Non-maintenance contract customers should proceed by using the enablers you have ordered.

Also, replace step 12 with the following updated step 12:

12. Enter the permanent authorization code once you receive the authorization code from the Sun Licensing Center:

   a. Start the `nwadmin` GUI.

   b. Click Server> Registration.

   c. Highlight the Update/7 enabler code.

d. Enter the authorization code.

   e. Click Apply.

⚠ **Caution –** The update enabler provides base functionality from the previous version. It does not add new functionality *nor* does it update any optional modules (Application Module for Oracle).

For more information, refer to Appendix C.

Updating a Solaris Master Server from Solstice Backup Release 4.2.6 to 6.1 (Solstice Backup Install Guide, page 44-47)

Ignore this section. Updating a Solaris Master Server from Solstice Backup Release 4.2.6 to 6.1 is *not* supported.

Server Software Products (Solstice Backup Admin. Guide, page 59)

Change the first sentence in the first bullet to: "Solstice Backup WorkGroup Edition enables you to back up the server and up to seven clients."

Autochanger Sharing Licensing and Availability (Solstice Backup Admin. Guide, page 379)

The Autochanger Sharing Licensing and Availability section states that the Network Edition of the Solstice Backup software requires an enabler license to use the autochanger sharing feature.

That is incorrect. Autochanger sharing is automatically enabled as part of the Network Edition of the Solstice Backup release 6.1 software.

Server Limitations for Three-Party Backups with NDMP (Solstice Backup Admin. Guide, page 450)

A note in the Three-Party Backups section states:

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.

This is *not* correct. Backups between an EMC NAS device and a NetApp server, or between a NetApp NAS device and an EMC server are *not* supported. The two NDMP servers must be of the same type. The following server limitations apply in a three-party backup:

- A NetApp server must back up and restore to another NetApp server.
- An EMC server must back up and restore to another EMC server.
- An Auspex server must back up and restore to another Auspex server.
- A SnapImage server must back up and restore to another SnapImage server.

Non-Destructive Save Set Recovery for NDMP Is Not Supported (Solstice Backup Admin. Guide, page 459)

The last paragraph in the section How to Perform a Destructive Save Set Recovery section states the following:

For additional information on save set recoveries, including recovery options, see "Save Set Recovery" on page 306.

This cross-reference is intended to provide general information about how a save set recovery works, as well as how to perform non-NDMP save set recoveries. Do *not* use the instructions described in the Save Set Recovery section for NDMP recoveries.

For NDMP recoveries, use the instructions provided in the section titled "Performing NDMP Recoveries" in Chapter 11: NDMP Support for the Solstice Backup Server.

Caution – The following save set recovery methods for non-NDMP recoveries are *not* supported for NDMP recoveries:

- Using the nwadmin GUI with anything other than the default value in the Paths to Recover field. The Paths to Recover field must contain the default path of the save set being recovered when performing an NDMP save set recovery, or the recovery fails.

NDMP Module

Network Data Management Protocol (NDMP) is a standards-based, storage management client/server protocol for enterprise-wide backup of heterogeneous, network-attached storage.

NDMP enables the Solstice Backup 6.0 and later software releases to provide connections to computers with NDMP data modules for tape operations, allowing a significant reduction in network traffic and better use of the time window dedicated to backups. NDMP also allows the Solstice Backup software to maintain the client index and media database, as well as control backup operations on a computer with an NDMP service installed.

NDMP Support

This section describes operating systems that support the NetWorker NDMP module, as well as the NDMP disk and tape services that are supported by the NetWorker NDMP module.

Operating Systems

The NetWorker NDMP module is supported on Solaris 2.6, 7 and 8.

NDMP Disk and Tape Services

The following NDMP disk and tape services are supported by the NetWorker NDMP module:

- NetApp Filer
- EMC Celerra, IP4700
- Auspex NS2000
- Legato Celestra®
- Legato NetWorker SnapImage Module

The Legato SnapImage Module is a high-performance backup solution that uses snapshot technology to back up and recover data across a LAN or SAN while keeping the filesystem available to users. The software provides full image backups and recoveries for customers with large filesystems who need to back up large amounts of data in a short period of time. The Legato SnapImage Module can be used as an NDMP tape server.

For additional information about SnapImage, refer to the *Legato SnapImage Installation and Administrator's Guide*.

Additional Features

Direct Access Restore (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 for large backups.

DAR is supported on NetApp, EMC Celerra, and Auspex. For information about which versions support DAR, refer to the Software Compatibility Guide on the Legato web site at `www.legato.com`.

When using DAR, NetApp cannot recover directories. Only files are recovered.

Important Notes and Tips

This section describes important notes and helpful tips related to the NetWorker NDMP module.

Updates to NDMP Documentation

Since the release of the Solstice Backup 6.1 documentation set, misleading documentation in the NDMP chapter of the *NetWorker Administrator's Guide, UNIX Version*, has been discovered.

Server Limitations for NDMP Three-Party Backups

For more information, see "Server Limitations for Three-Party Backups with NDMP (Solstice Backup Admin. Guide, page 450)" on page 62.

Non-Destructive Save Set Recovery for NDMP Not Supported

For more information, see "Non-Destructive Save Set Recovery for NDMP Is Not Supported (Solstice Backup Admin. Guide, page 459)" on page 63.

Licensing Update

Prior to the release of Solstice Backup 6.1 software, each NDMP client required both an NDMP client connection and a SAN storage node enabler code. Solstice Backup 6.1 software does not require the SAN storage node enabler code. Only an NDMP client connection is required.

Environment Variable Updates

The environment variable values specified in Chapter 12 of the *Solstice Backup 6.1 Administrator's Guide* 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 and syntax, as they may change periodically.

Required NDMP Version When Running DAR on NetApp DataONTAP 6.1

To run DAR on a NetApp system, you must have NDMP Version 3 and Solstice Backup 6.1 or later software. Recoveries fail if you use an earlier NDMP version or an earlier Solstice Backup software release.

To verify the NDMP version you have:

1. Log in to the NetApp DataONTAP 6.1 system.

2. At the prompt, type **ndmpd status**.

   The NDMP status is displayed. If an NDMP process is running, the status of that process is displayed. If an NDMP process is not running, whether ndmpd is on or off is displayed.

To change the NDMP version:

1. At the prompt, type **ndmpd off** to stop the NDMP process.

   The NetApp host confirms that the NDMP process has stopped.

2. Type **ndmpd** version 3 to change the NDMP version.

   The NetApp host confirms that the NDMP version is now set to 3.

3. Type **ndmpd** on to restart the NDMP process.

The NetApp host confirms that the NDMP process has been started. You can now run DAR.

Required Variable for Restoring ACLs when Using DAR with a NetApp Host

To restore ACLs when using DAR with a NetApp host, you must set the environmental variable EXTRACT_ACL=T. When using DAR, the default value is F for false. You must set the variable to T for true. The variable is case-sensitive.

Required Variables for Using DAR with an EMC Celerra Host

Table 2-1 lists the required environment variables when using DAR with an EMC Celerra host. All variables are case-sensitive.

TABLE 2-1   Required Variables for Using DAR with an EMC Celerra Host

Environment Variable	Definition
DIRECT=y/n	y/n. This value must be set to y when using DAR. If you are not performing a DAR recovery, the value is set to n. For additional information about DAR, see "Direct Access Restores" in the *Legato NetWorker Administrator's Guide*.
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.
EXTRACT_ACL=T	To restore ACLs when using DAR with an EMC Celerra host, you must set the value to T.

For additional environmental variables used with an EMC Celerra server, refer to the *Solstice Backup 6.1 Administrator's Guide*.

Backing up Bootstrap Information during a Command Line Backup

When you perform a backup through the NetWorker Administrator interface, the bootstrap information is updated automatically. However, if you use the nsrndmp_save command to perform the backup, the bootstrap information is not automatically updated.

To make sure the bootstrap information is updated automatically, you must include the -s *save_group* option in the backup command. The -s *save_group* option specifies the group that needs to be backed up.

Solstice Backup Server Configuration Requirements for NDMP Support

When performing Solstice Backup operations using NDMP, additional memory, swap space, and `/nsr/tmp` space are required. The following sections describe these requirements.

Memory and Swap Space Requirements

Solstice Backup backup and recovery operations using NDMP require additional memory and/or swap space. Use the following guidelines when performing a backup or recovery operation:

■ For a backup, you need 204 bytes of memory for each directory or file entry to be backed up, plus an additional 10 KB of memory and/or swap space per backup session. For example, backing up one million files requires 204 MB of memory/swap space, plus 10 KB for the backup session, which equals 204.01 MB total memory and/or swap space. See the section below for information about the additional `/nsr/tmp` space requirements

■ Solstice Backup browsable recoveries (non-save set recoveries) also require additional memory and/or swap space, approximately 450 bytes for each directory or file entry to be recovered. For example, recovering one million files requires 450 MB of memory and/or swap space.

Note – In a 32-bit operating system (on either a 32- or 64-bit architecture), the maximum memory that can be allocated to a process is 2 GB. Therefore, recovering more than 4 million files in a single operation is not possible. To recover more than 4 million files, use either a series of recovery operations with a 32-bit operating system or appropriate memory and/or swap space with a 64-bit operating system.

If the Solstice Backup server does *not* have enough swap space to hold the entire file history, the following error message appears:

```
cannot decode the file
```

Space Requirements for `/nsr/tmp`

The space requirements for the `/nsr/tmp` directory are as follows:

- Solstice Backup backup operations using NDMP require temporary work space in `/nsr/tmp`. Approximately 260 bytes of disk space must be available for each directory or file entry. For example, to back up one million files requires 260 MB of `/nsr/tmp` space.

- Solstice Backup recovery operations do *not* require `/nsr/tmp` disk space.

Optimizing Qtree Backups with NetApp Data ONTAP 5.3.x

If you are using Data ONTAP 5.3.*x* and are backing up more than eight qtrees on a single 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, type the following 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 your Data ONTAP documentation.

Environmental Variable Considerations

The environmental variables to be set for NDMP backups are case sensitive. You must match the case that is listed in Chapter 12: NDMP Support for the Solstice Backup Server in the *Solstice Backup* 6.1 *Administrator's Guide*.

For example, for HIST=y, HIST must be all uppercase letters and y must be a lowercase letter. If you do not use the correct case when specifying a value for the variable, the default value is used.

All Hosts for Directed Recoveries Must Be Solstice Backup Clients

When performing a directed recovery to a host other than the original host, the host to which data is being recovered must be a client of the Solstice Backup server. If it is *not* a client, the server cannot obtain the correct user name and password, and the data cannot be recovered to that host.

Installation and Upgrade Notes

This section includes topics relevant to installing and upgrading Solstice Backup software in an NDMP environment.

Enabling Dynamic Drive Sharing in an NDMP Environment

All the components of a SAN configuration must be compatible when Dynamic Drive Sharing (DDS) is enabled with the NetWorker NDMP module:

- The fiberchannel switches must be compatible with the NDMP filers within a SAN.
- NDMP filers and libraries in the SAN must also be compatible with each other.

Drives can be shared among NDMP filers in a SAN environment, but the NDMP devices must be 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
- EMC IP4700 to EMC IP4700
- NetApp to NetApp (any NetApp filers that Network Appliance supports within a zone)

For further information regarding compatibility issues, refer to the following:

- Legato Compatibility Guides at `www.legato.com`
- NDMP vendors' web sites
- Tape library vendors' web sites

For a brief description of the DDS feature, see .

Known Limitations

This section describes known limitations, configuration tips, and workarounds in NDMP module support.

User Must Be Root for NDMP Backup or Recovery

In order to perform either a backup or recovery operation, the user must be logged in as root on the Solstice Backup server. If the user is not root, the backup or recovery fails.

Permission Problems with Index Recoveries of EMC Celerra Servers (Legato Bug ID: LGTpa29103)

When using the `nwrecover` GUI or the `recover` command to perform an index-based recovery of one or more files to a directory that does not exist in the target EMC Celerra server, EMC Celerra creates a new directory. However, the attributes (permissions and ownership) of this new directory will not match the attributes of the originally backed up directory.

To avoid this problem, perform one of the following:

- Manually change the directory attributes of the new directory after the recovery.
- Perform a save set recovery instead of an index-based recovery.

Large Number of Files Limitation with Index Recoveries on EMC Celerra

When using the `nwrecover` GUI or the `recover` command to perform an index-based recovery of an EMC Celerra server, the recovery is recommended *not* to exceed 100,000 files per recovery.

To recover a larger number of files, perform a save set recovery through the `nwadmin` GUI or the `nsrndmp_recover` command.

Disabling Locally-Attached Drives Can Cause Restore Failure (Legato Bug ID: LGTpa31403)

If you have library sharing between two data movers and you use the `nwadmin` GUI to disable one or more drives that are locally attached to one of those data movers, performing a restore forces NetWorker to choose the drive that is attached to the other data mover, thus causing a three party restore. This forced three party restore then fails. To successfully restore the data, re-enable the locally-attached drives.

NDMP Save Set Recovery Fails with Files and Directories with Long Names (Legato Bug ID: LGTpa31404)

If a file or directory name contains approximately 1 KB of characters or longer, recovering these files through an NDMP Save Set recovery fails. To avoid this problem, create files and directories with shorter names.

NDMP Index Recovery Fails with Certain File and Directory Names (Legato Bug ID: LGTpa31405)

An NDMP index recovery might fail if the file or directory names contain the following characteristics:

- Approximately 1 KB of characters or longer
- Specific metadata characters

To avoid these problems, create files or directories with shorter names and do *not* include metadata characters.

No File Index for Multiple Mount Points on the Same Host (Legato Bug ID: LGTpa24557)

In this release, the Solstice Backup software does *not* distinguish between mount points and normal directories when sending file history to `nsrndmpix`. For example, if you back up a filesystem mounted on `/mnt2/mnt3` and browse the file indexes for the filesystem, the software displays all the indexes correctly. If you then unmount `/mnt2/mnt3` and later back up only `/mnt2`, nwrecover and the `recover` command do *not* show the directories under `/mnt2/mnt3`. In order to see the indexes from `/mnt2/mnt3`, you must browse the file indexes of the previously backed-up partition.

You can avoid this issue by backing up `/mnt2` first and then backing up `/mnt2/mnt3`. When the backups are completed in this order, the software displays all the indexes correctly.

NDMP Backup Pools (Legato Bug ID: LGTpa23888)

Recovery of an NDMP backup fails if the volume is loaded in a non-NDMP device. To prevent non-NDMP volumes from being recommended for NDMP backups, create separate pools for NDMP and non-NDMP volumes.

File Recovery Limitation (Legato Bug ID: LGTpa24935)

When performing a recovery while also using Data ONTAP 6.0, you must specify a forward slash (/) at the end of the mount point. Otherwise, the recovery fails. When you specify a forward slash, the recovery is successful.

Pre-Initialized STK Tapes in NDMP-Enabled Tape Devices Not Used (Legato Bug ID: LGTpa28778)

STK offers two types of tapes: pre-initialized and uninitialized. Solstice Backup software does not use tapes that have been pre-initialized in an NDMP-enabled tape device.

To ensure that all tapes in an NDMP tape device are used:

- Do *not* use tapes that were pre-initialized in an NDMP-enabled tape device.
- Label pre-initialized tapes in a non-NDMP tape device using Solstice Backup software and then use only those tapes in the NDMP-enabled tape devices.

Path to Recover Attribute in nwadmin Not Supported for NDMP (Legato Bug ID: LGTpa27162)

When performing a recovery using the nwadmin program, you must complete the attributes in the Save Set Recover Status window. If, however, you are performing an NDMP recovery, do *not* complete the Path to Recover attribute in the Save Set Recover Status window. This attribute is *not* supported in an NDMP recovery; attempting to use it results in an error message. Leave the Path to Recover attribute blank and complete the other required attributes before clicking Start.

EMC Celerra Level 5 Recovery Recovers More than Original Backup (Legato Bug ID: LGTpa29467)

During a level 5 recovery on EMC Celerra, the Solstice Backup server also recovers an empty subdirectory from the last full backup, even though the contents and name of the subdirectory did not change. To avoid this problem, perform a file-by-file recovery when recovering a level 5 backup.

NDMP Recovery of Large Number of Files Causes NetApp Filer Panic (Legato Bug ID: LGTpa32106)

If you perform a file-by-file recovery using the `recover` command or the `nwrecover` GUI, specifying a recovery of more than 10,240 files may cause the NetApp Filer to panic or result in an error message.

To avoid this problem, perform a save set recovery or recover only 10,240 files or less at one time.

Network Appliance has documented this problem in NetApp Bug 52946 and has patches available for both Data ONTAP 6.1 and 6.1.1. These Data ONTAP patches are available by contacting Network Appliance Customer support at *www.netapp.com*.

Specify Mount Point Field when Performing a NDMP Directed Recovery (Legato Bug ID: LGTpa29497)

When you perform a directed recovery on an EMC Celerra or NetApp client, the recovery fails if a value has been specified in the Relocate field of the Options window.

To avoid this problem and perform a directed recovery, specify the new location in the Mount Point field of the Save Set Recover Status window and leave the Relocate field blank.

Only New or Changed Files Appear in GUI for NDMP Backups (Legato Bug ID: LGTpa28760)

When either the `nwrecover` or `recover` command is run after a level 1 through 9 backup, both commands show only the files saved during that backup, rather than all the files from the last full backup.

To view all the files from the most recent full backup, change the browse time to a time before the level 1 through 9 backup was performed.

The jbconfig Program Configures Drive with Wrong Hostname When Enabling DDS on NDMP Drive (Legato Bug ID: LGTpa30508)

When configuring a remote NDMP device using the `jbconfig` program, the program might configure the device with the wrong hostname if all following are true:

- The `jbconfig` program is being run from a storage node.
- The drive is local to the Solstice Backup server.
- The drive has not been created prior to running the `jbconfig` program.

The result is either an incorrectly configured system or failed configuration.

To share an NDMP device successfully under these conditions:

1. Create the device through the `nwadmin` GUI:

2. From the Media menu, select Devices.

3. Click Create.

4. Create the device in the following format:

```
rd= server_name:/device_path
```

NDMP Backup Hangs after Data Transfer (Legato Bug ID: LGTpa30208) (Resolved in Solstice Backup 6.1 with Patch 111962-xx)

When an NDMP backup is performed in a Solaris environment, the backup might hang after the data is transferred. The data is successfully backed up and the backup process finishes, but there could be a delay between the data transfer and the completion of the backup. For example, if a backup took two hours to transfer the data, the process could hang for an additional few hours.

NDMP Recovery Fixes Combined and Available for Solstice Backup 6.1 (Legato Bug IDs: LGTpa30748, 31270) (Resolved in Solstice Backup 6.1 with Patch 111962-xx)

For more information see,"NDMP Recovery Fixes Combined and Available for Solstice Backup 6.1 (Legato Bug IDs: LGTpa30748, 31270) (Resolved in Solstice Backup 6.1 with Patch 111962-xx)" on page 24.

Installing Solstice Backup in a Sun Cluster 3.0 Environment

This chapter explains how to install Solstice Backup software in a Sun® Cluster 3.0 environment. It includes the following sections:

A Sun Cluster 3.0 system enables applications to be run and administered as highly available resources (data services). For detailed information about Sun Cluster 3.0, refer to the *Sun Cluster System Administration Guide*.

Caution – Highly available storage nodes are *not* supported.

Cluster Terminology and Concepts

This section discusses the following topics:

-
-
-

Cluster Terminology

For consistency and clarity, this section defines the following terms as follows:

Node — a physical machine that is a member in the cluster.

Failover node — a cluster node to which the highly available Solstice Backup server can fail over.

Virtual client — a Solstice Backup client that is associated with a highly available or scalable service in a cluster.

Cluster client — a Solstice Backup cluster client is either a physical client or a virtual client. A physical client resource backs up data belonging to a physical node within a cluster. A *virtual* client resource backs up data belonging to a highly available or scalable service within a cluster.

Stand-alone Solstice Backup server — a Solstice Backup server that is *not* highly available.

Cluster-Aware Solstice Backup server — a Solstice Backup server that can recognize cluster clients.

About Clustering

Clustering means that two or more nodes are connected and appear to network users as a single, highly available system. A highly available system allows the application services to continue despite most hardware or software failures. Each *cluster node* has its own IP address. Each cluster node might also have *private* (local) resources or disks that are available only to that computer. In addition to the nodes, a cluster can have one or more *highly available services*, each with its own IP address. A *highly available service* can have zero or more shared disks.
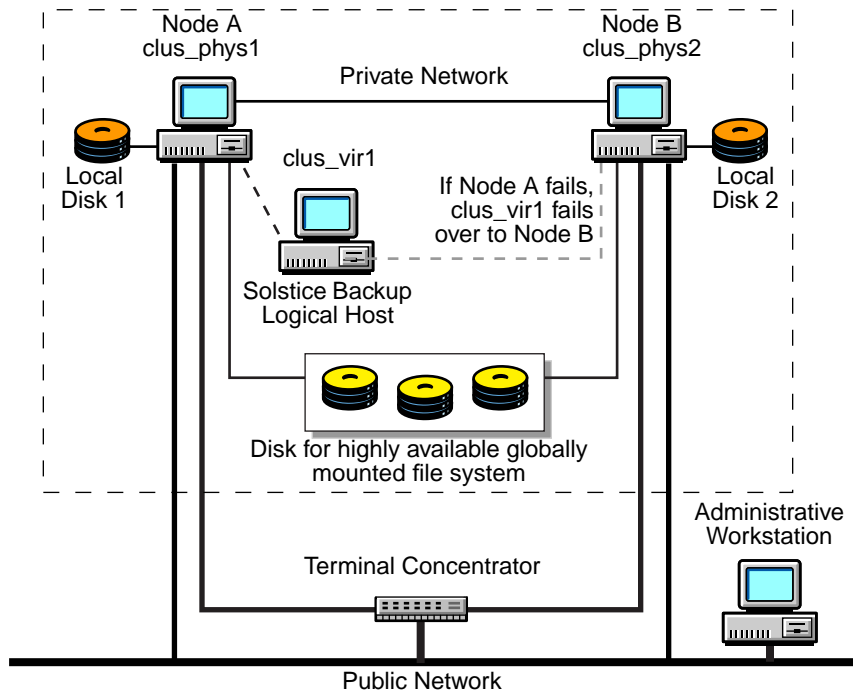
FIGURE 3-1　General Cluster Configuration

displays a general cluster configuration consisting of two or more nodes and at least one Solstice Backup server. In this illustration:

- Node A is a computer with its own **Local Disk 1**.
- Node B is a computer with its own **Local Disk 2**.
- Highly available Solstice Backup Server owns **Shared Disks**.

In this example, the Solstice Backup server can fail over between Node A and Node B; however, the Solstice Backup server only runs on one node at a time.

The Solstice Backup server resource group contains various resources including a logical hostname resource. Although the Solstice Backup server software is installed on all the physical nodes within the cluster, the server runs on only one node at a time. If a hardware or software breakdown occurs on the node on which the Solstice Backup server is running, it moves to another node in the cluster.

The Solstice Backup client software runs on all the physical nodes within the cluster, regardless of where or whether-or-not the Solstice Backup server resource group is running. Only one instance of the client software (nsrexecd) runs on each physical node within the cluster. The Solstice Backup client software is designed to recognize more than one client (physical client plus virtual clients) that might be associated with a single physical node.

Software Installation Roadmap

The following tasks are a roadmap to follow when installing the Solstice Backup software:

1. Read this roadmap and the referenced sections before installing the software.

2. Review the requirements section and note the default directory location. For more information, see "Installation Requirements" on page 83.

3. Install the required Solstice Backup software. For more detailed information, see "Installing Solstice Backup Software" on page 94.

4. Register and license all of your Solstice Backup products.

Installation Requirements

The following sections discuss the software and hardware required to install and configure the Solstice Backup server or client software within a Sun Cluster 3.0 environment:

For the most up-to-date information regarding software and hardware requirements, refer to the Legato Compatibility Guides on the Legato web site at `www.legato.com`.

Software Requirements

To install a highly available Solstice Backup server, make sure that the following software is installed on each node in the cluster:

- Solaris 8 operating environment
- Sun Cluster version 3.0 and update releases
- Volume Manager software (Solstice Disksuite or Veritas Volume Manager)
- Solstice Backup release 6.1 with patch 111962-xx

Note – With Solstice Backup release 6.1 L10N software, patch 111962-xx is *not* required.This problem only occurs with EMC Celerra version 2.2.xx.x.

Hardware Requirements

To install the Solstice Backup software, make sure that the following hardware requirements are met for server installation only:

- Multi hosted disks that are used as a mount point for global file systems to contain the shared `/nsr` area.
- A device with local affinity for the local bootstrap backup that is connected to all the nodes within the cluster.

Configuration Options

See "Solstice Backup for Sun Cluster 3.0 Administration" on page 125 for information on how to configure:

- Tape libraries with a highly available Solstice Backup server
- Remote physical storage node with a highly available Solstice Backup server

System Information Requirements for a Highly Available Solstice Backup Server

FIGURE 3-2 and TABLE 3-1 display the system information needed to install a highly available Solstice Backup server within a Sun Cluster 3.0 environment.
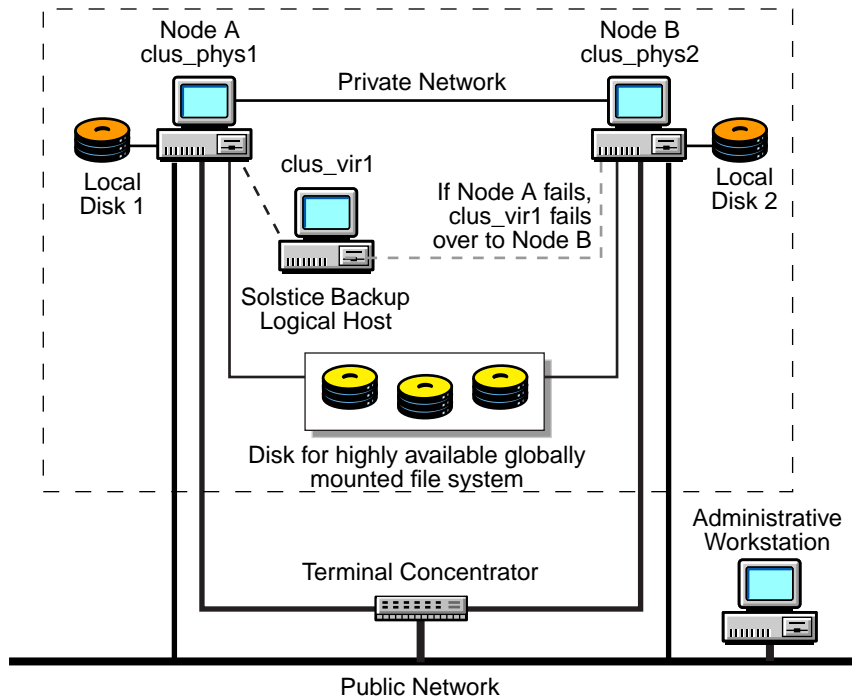
FIGURE 3-2   Sample Cluster Configuration

TABLE 3-1   Required Information: Highly Available Solstice Backup Server

Information Required	Example
Solstice Backup logical hostname with an IP address in the same subnet as the cluster nodes	`clus_vir1` `192.168.1.10`
Solstice Backup globally mounted configuration area	`/global/nw`
Solstice Backup cluster configuration script	`/usr/sbin/nsr/networker.cluster`
Solstice Backup start script	`/usr/sbin/nsr/networker.start`
Solstice Backup stop script	`/usr/sbin/nsr/networker.stop`
Solstice Backup monitor start script	`/usr/sbin/nsr/networker.fm_start`
Solstice Backup monitor stop script	`/usr/sbin/nsr/networker.fm_stop`
File containing the hostids of all the cluster members	`/nsr/res/hostids`

Updating from a Previous Release of Solstice Backup Software

Use the following procedures to update or upgrade from previous releases of Solstice Backup software to Solstice Backup release 6.1:

- "To Update from Solstice Backup Release 6.0 on Sun Cluster 2.2" on page 85
- "To Upgrade from a Standalone Solstice Backup Environment" on page 88

▼ To Update from Solstice Backup Release 6.0 on Sun Cluster 2.2

Use the following procedure to update from Solstice Backup release 6.0 on Sun Cluster 2.2 to Solstice Backup release 6.1 on Sun Cluster 3.0.

Note – In order to update to Solstice Backup release 6.1, you must also update to Sun Cluster 3.0 in a Solaris 8 operating environment.

To update from Solstice Backup release 6.0 on Sun Cluster 2.2:

1. Make sure that your PATH environment variable includes /usr/sbin and /opt/sunwcluster/bin.

2. On the node currently mastering the Solstice Backup logical host:

   a. Make sure that all of your backups are up-to-date.

   b. Shut down the Solstice Backup software. For example:

```
# hareg -n networker
# /usr/sbin/nsr/nsr_shutdown
```

   c. Save a copy of the server's /nsr databases. These global databases are located in the shared disks within the cluster and will be needed after the upgrade.

   You can use *one* of the following methods:

   – Save the global databases to a file that you can later ftp to a remote computer for safekeeping:

```
# cd /logical_hostname/nsr
# /usr/lib/nsr/uasm -s -i . > filename
```

Caution – Make sure that you do not overwrite the Solstice Backup storage tapes.

   – Save the global databases to a local tape drive:

```
# cd /logical_hostname/nsr
# /usr/lib/nsr/uasm -s -i . | dd of=tapename
```

   – Save the global databases directly to a remote file or tape drive:

```
# cd /logical_hostname/nsr
# /usr/lib/nsr/uasm -s -i . | rsh remote_machine dd of=remote_filename
```

3. For all nodes in the cluster:

   a. Install the Solaris 8 operating environment and Sun Cluster 3.0 software.

   b. Boot the computer in cluster mode.

   For detailed information about how to install Solaris 8 operating environment and Sun Cluster 3.0 software, refer to appropriate Sun documentation.

4. On each node in the cluster, install the Solstice Backup software. Perform the following tasks:

   - "Task 1: Install the Solstice Backup Software in a Cluster" on page 95
   - "Task 2: Define the Solstice Backup Server as Highly Available" on page 98
   - "Task 3: Create an Instance of the Solstice Backup Server Resource Group" on page 99

5. On the node that is currently running the Solstice Backup resource group:

   a. Stop the Solstice Backup resource group:

```
# scswitch -F -g resourcegroup_name
```

   b. Go to the global /nsr directory and delete its contents:

```
# cd /global/nw/nsr
# rm -rf *
```

   c. Restore the contents of the /nsr directory to the global area.

   You can use *one* of the following methods:

   – Copy the file back to the server and restore it:

```
# /usr/lib/nsr/uasm -r < filename
```

   – Recover the file from the local tape drive:

```
# dd if=tapename | /usr/lib/nsr/uasm -r
```

   – Recover the file from the remote file or tape:

```
# rsh remote_machine dd if=remote_filename | /usr/lib/nsr/uasm -r
```

6. Restart the Solstice Backup software:

```
# scswitch -Z -g resourcegroup_name
```

7. To continue configuring the Solstice Backup server perform the following tasks:

Note – You do *not* require a new update enabler and authorization key. Additional registration and licensing procedures are *not* required.

Caution – Do not change the logical hostname for the highly available Solstice Backup server. If you change the logical host configuration for the highly-available Solstice Backup server after you update the software, you must permanently license and authorize the highly available Solstice Backup server.

▼ To Upgrade from a Standalone Solstice Backup Environment

You can upgrade Solstice Backup from following non-cluster releases to Solstice Backup release 6.1 on Sun Cluster 3.0:

- Solstice Backup release 6.0 to 6.1
- Solstice Backup release 5.5.1 to 6.1
- Solstice Backup release 5.1/5.1.1 to 6.1

Caution – Upgrading from a non-cluster Solstice Backup server to a highly available Solstice Backup server is *only* supported if the physical hostname of the pre-existing server is defined as a logical hostname of the newly defined cluster. This logical hostname *must* be used by the highly available Solstice Backup service.

To upgrade Solstice Backup from non-cluster releases:

1. On the standalone Solstice Backup server, update to Solstice Backup release 6.1. For detailed instructions, refer to the *Solstice Backup 6.1 Installation Guide*.

2. On Solstice Backup server:

   a. Make sure that all of your backups are up-to-date.

   b. Shut down the Solstice Backup software. For example:

```
# /usr/sbin/nsr/nsr_shutdown
```

   c. Save a copy of the server's /nsr databases. Use *one* of the following methods:

      – Save the databases to a file that you can later ftp to a remote computer for safekeeping:

```
# cd /nsr
# /usr/lib/nsr/uasm -s -i . > filename
```

⚠ **Caution –** Make sure that you do not overwrite the Solstice Backup storage tapes.

      – Save the databases to a local tape drive:

```
# cd /nsr
# /usr/lib/nsr/uasm -s -i . | dd of=tapename
```

      – Save the databases directly to a remote file or tape drive:

```
# cd /nsr
# /usr/lib/nsr/uasm -s -i . | rsh remote_machine dd of=remote_filename
```

3. Upgrade the following software and incorporate the node into a cluster:

   a. Install and configure the Solaris 8 operating environment software.

   b. Install and configure the Sun Cluster 3.0 software.

Note – Make sure that you change the physical hostname of the node.

---

   c. Boot the computer in cluster mode.

   For detailed information about how to install Solaris 8.0 operating environment and Sun Cluster 3.0 software, refer to appropriate Sun documentation.

4. On each node in the cluster, install the Solstice Backup software. Perform the following tasks:

   - "Task 1: Install the Solstice Backup Software in a Cluster" on page 95 for each node that does not already have the Solstice Backup release 6.1 software installed.

   - "Task 2: Define the Solstice Backup Server as Highly Available" on page 98

   - "Task 3: Create an Instance of the Solstice Backup Server Resource Group" on page 99

5. On the node that is currently running the Solstice Backup resource group:

   a. Stop the Solstice Backup resource group:

   ```
   # scswitch -F -g resourcegroup_name
   ```

   b. Go to the global /nsr directory and delete its contents:

   ```
   # cd /global/nw/nsr
   # rm -rf *
   ```

   c. Restore the contents of the /nsr directory to the global area.

      You can use *one* of the following methods:

      – Copy the file back to the server and restore it:

   ```
   # /usr/lib/nsr/uasm -r < filename
   ```

– Recover the file from the local tape drive:

```
# dd if=tapename | /usr/lib/nsr/uasm -r
```

– Recover the file from the remote file or tape:

```
# rsh remote_machine dd if=remote_filename | /usr/lib/nsr/uasm -r
```

6. Restart the Solstice Backup software:

```
# scswitch -Z -g resourcegroup_name
```

7. Reconfigure the Solstice Backup server to reflect the new cluster configuration. For example, you might need to redefine the storage devices.

   For detailed instructions, see "Using Tape Libraries and Stand-Alone Tape Devices" on page 136.

8. Continue configuring the Solstice Backup server. Perform the following tasks:

   ■ "Task 4: Grant Access to the Highly Available Solstice Backup Server" on page 101
   ■ "Task 5: Configure the Solstice Backup Server" on page 104
   ■ "Task 6: Configure Clients under the Solstice Backup Server" on page 104
   ■ "Task 7: Create Instances of the Solstice Backup Client Resource Type" on page 106

9. Register the software. For details see, "Register and License the Highly Available Solstice Backup Server" on page 92.

Register and License the Highly Available Solstice Backup Server

The registration and licensing requirements differ depending upon the version of Solstice Backup software you are upgrading from:

> **Caution –** Network Edition does *not* support backup and recover for a highly available Solstice Backup server. A highly available Solstice Backup server requires a Power Edition enabler and the appropriate number of Cluster Client Connection licenses.

Registering Solstice Backup Release 6.0 to 6.1

> **Note –** A new update enabler is *not* required.

To register the software, perform the following tasks:

Registering Solstice Backup Release 5.5.1 to 6.1

> **Note –** A new Update/7 enabler *is* required. Only customers with maintenance contacts will automatically receive Update/7 enablers.

To register the software:

1. If you do *not* have a maintenance contract, order appropriate licenses for the highly available Solstice Backup server and other desired features.

2. Perform the following tasks:

Registering Solstice Backup Release 5.1/5.1.1 to 6.1

Note – A new Update/7 enabler *is* required. Only customers with maintenance contacts will automatically receive Update/7 enablers.

To register the software:

1. If you do *not* have a maintenance contract, order appropriate licenses for the highly available Solstice Backup server and other desired features.

2. Perform the following tasks:

   - "Task 8: Register and License the Highly Available Solstice Backup Server" on page 108
   - "Task 9: Create a Composite Hostid" on page 109
   - "Task 10: Authorize the Software" on page 110

Installing Solstice Backup Software

To install the Solstice Backup software, see the following sections:

▼ To Install a Highly Available Solstice Backup Server

To install and configure the Solstice Backup software, perform these steps:

Task 1: Install the Solstice Backup Software in a Cluster

The Solstice Backup software must be installed on all nodes in the cluster that you want to run the Solstice Backup resource group.

Note – Before you begin, make sure that the Solaris 8 operating environment and Sun Cluster 3.0 software is already installed on all cluster nodes and that those nodes boot in the cluster mode.

Also, make sure that your PATH environment variable includes /usr/sbin and /usr/cluster/bin.

To install the Solstice Backup software on nodes that are currently running the Solstice Backup resource group:

1. Access the Solstice Backup software from the distribution media. For detailed instructions, refer to the *Solstice Backup 6.1 Installation Guide*.

2. Keep a copy of your current configuration. The Solstice Backup software installation script modifies the /etc/rpc and /etc/syslog.conf files during the installation process. For example:

```
# cp /etc/rpc /etc/rpc.old
# cp /etc/syslog.conf /etc/syslog.conf.old
```

3. Enter the **pkgadd** command at the system prompt. For example:

```
# /usr/bin/pkgadd -d /cdrom/cdrom0/solaris/sparc
```

The following Solstice Backup software packages are listed:

1. `SUNWsbuc` (client software package)
2. `SUNWsbum` (optional man pages)
3. `SUNWsbun` (storage node software package)
4. `SUNWsbus1` (server software package)
5. `SUNWsbus2` (device drivers package)

4. The Solstice Backup software packages depend on each other. At the prompt, type the following:

1, 5, 3, 4, 2

This ensures that the packages are installed in the correct order.

5. Depending on which version of the Solstice Backup software you are installing, the installation steps differ:

- If you are installing Solstice Backup release 6.1 L10N software:

  a. When prompted to start the Solstice Backup daemons, type **yes**.

  b. When the prompt appears after the Solstice Backup daemons have started, type **q** to exit.

- If you are installing the Solstice Backup release 6.1 with patch 111962-xx:

  a. When prompted to start the Solstice Backup daemons, type **no**.

  b. When you have added all the applicable packages, and the prompt appears type **q** to exit.

  c. Download and install patch 111962-xx from the Sun web site. For detailed information see, "Downloading and Installing Patch 111962-xx:" on page 97.

Downloading and Installing Patch 111962-xx:

To download and install patch 111962-xx from the Sun web site:

1. In a web browser, enter the following address:

```
http://www.sun.com/service
```

2. Click on the Patches link at the *Online Service Center* section and navigate to the *Sunsolve Online* page.

3. Click on the Patchfinder link at the *Sunsolve Contents* section and navigate to the *Sunsolve Online Patches* page.

4. Enter the patch ID 111962 in the text field entry box, click Find Patch, and navigate to the *README* page for the patch.

Note – Patchfinder can find a patch with only its patch ID. The revision number -xx is not required.

---

5. Click on either the HTTP or FTP link at the top of the page to download the patch. The downloaded patch is in compressed tar format.

6. Uncompress and untar the patch:

```
# cd download_path
# /bin/uncompress 111962-xx.tar.Z
# /bin/tar xvf 111962-xx.tar
```

7. Enter the patchadd command at the system prompt. For example:

```
# cd download_path
# /usr/sbin/patchadd 111962-xx
```

8. Verify that the Solstice Backup daemons are running. For example:

```
# ps -ef | grep nsr
```

9. If the daemons are *not* running, enter the following command to start them.

```
# /etc/init.d/networker start
```

Task 2: Define the Solstice Backup Server as Highly Available

To define and configure the Solstice Backup server as highly available:

1. Log in as root.

2. Make sure that the /etc/hosts file on each cluster node contains the name of the logical host. You can also publish the logical hostname in the Domain Name System (DNS) or Network Information Service (NIS).

3. Make sure that a globally mounted file system is available for the highly available Solstice Backup server.

4. From each node in the cluster that will run the Solstice Backup server process:

   a. Run the cluster configuration script, /usr/sbin/nsr/networker.cluster. For an example of the script, see "Solstice Backup Cluster Script" on page 122.

      This script defines the LGTO.serv and the LGTO.clnt resource types that the Solstice Backup software requires.

   b. In response to the following prompts, enter the following information:

   ```
   Enter directory where local NetWorker database is installed
   [/nsr]?
   ```

   – Enter the location of the local Solstice Backup database directory provided during the installation procedure. For example: /space/nsr

   ```
   Do you wish to configure for both NetWorker server and client?
   Yes or No [Yes]?
   ```

   – Enter **yes** to configure the server software. This will also install the client software by default.

   – Entering **no** configures only the client software.

   ```
   Do you wish to add now the site-specific values for:
   NSR_SHARED_DISK_DIR and NSR_SERVICE_ID
   Yes or No [Yes]?
   ```

– Enter **yes**. This ensures compatibility with other cluster environments.

```
Enter the Logical Hostname to be used for NetWorker?
```

– Enter the published logical hostname to be used by the highly available Solstice Backup server. For example: *clus_vir1*

```
In what path will this nsr directory be created/located?
```

– Enter the pathname of the globally mounted nsr directory that will contain the configuration information for the highly available Solstice Backup server. For example: /global/nw.

Note – The examples used in this procedure are in reference to the examples found in "Required Information: Highly Available Solstice Backup Server" on page 85.

If you make mistakes during this configuration task, you can undo any changes to the configuration by running the networker.cluster -r option and then running the networker.cluster script again.

Task 3: Create an Instance of the Solstice Backup Server Resource Group

A resource group must be created for the highly available Solstice Backup server. It must contain the following resources:

- Logical hostname
- Globally mounted file system for the /nsr area
- LGTO.serv resource
- LGTO.clnt resource

To create an instance of the Solstice Backup server resource group, perform the following steps on one node in the cluster:

1. Create a resource group:

```
# scrgadm -a -g resourcegroup_name
```

2. Add the logical host:

```
# scrgadm -a -L -g resourcegroup_name -l clus_vir1
```

3. Create an instance of the LGTO.serv resource type that has a dependency on the logical host. You must also specify a configuration directory for the global /nsr area. For example:

```
# scrgadm -a -j server -t LGTO.serv -g resourcegroup_name \
-y resource_dependencies=clus_vir1 -x config_dir=/global/nw
```

4. Create an instance of the LGTO.clnt resource type. For example:

```
# scrgadm -a -j client -t LGTO.clnt -g resourcegroup_name \
-x clientname=clus_vir1 -x owned_paths=/global/nw
```

5. Start the Solstice Backup resources. For example:

```
# scswitch -Z -g resourcegroup_name
```

Example: A Highly Available Solstice Backup Server

In the following example, the highly available Solstice Backup server uses the logical hostname *backup_server*. The highly available Solstice Backup server uses */global/sbu* (globally mounted file system) as its configuration directory.

1. Create a resource group with the name *backups.*

```
# scrgadm -a -g backups
```

2. Add the logical hostname resource type to the resource group created in the previous step.

```
# scrgadm -a -L -g backups -l backup_server
```

3. Create an instance of the `LGTO.serv` resource type with the name *sbu_server*. This resource belongs to the resource group *backups* and has a dependency on the logical host created in the previous step. You must also specify the configuration directory on the globally mounted file system */global/sbu*.

```
# scrgadm -a -j sbu_server -t LGTO.serv -g backups \
-y resource_dependencies=backup_server -x config_dir=/global/sbu
```

4. The Solstice Backup logical host is also a client of the highly available Solstice Backup server. You must create an instance of the `LGTO.clnt` resource type for the logical host *backup_server* within the resource group *backups*. The name of this resource is *sbu_client*.

```
# scrgadm -a -j sbu_client -t LGTO.clnt -g backups \
-x clientname=backup_server -x owned_paths=/global/sbu
```

5. Start the highly available service associated with the resource group *backups*.

```
# scswitch -Z -g backups
```

Task 4: Grant Access to the Highly Available Solstice Backup Server

Before a Solstice Backup server can back up a client, the client must grant the server access. Granting access is controlled by the `servers` file. For a node running the highly available Solstice Backup server, this file is located in the globally mounted file system. Otherwise, this file is located on a local disk.

If the `servers` file does *not* exist, any Solstice Backup server can back up this client. If the `servers` file *does* exist, only the specific servers listed in the file can back up the client.

To grant access to the highly available Solstice Backup server:

1. On one node in the cluster:

   a. Stop the Solstice Backup service using the cluster management software. For example:

   ```
   # scswitch -F -g resourcegroup_name
   ```

   b. Edit or create the servers file in the globally mounted /nsr area. For example /global/nw/res/servers:

   – Add the set of Solstice Backup servers, one per line, that you want to grant access to this client.

   – Add an entry for the Solstice Backup logical hostname first. Then add entries for each physical host that can run the Solstice Backup resource group. For example:

   clus_vir1

   clus_phys1

   clus_phys2

2. On each node in the cluster:

   a. Shut down the Solstice Backup processes and verify that all Solstice Backup daemons have stopped. For example:

   ```
   # /usr/sbin/nsr/nsr_shutdown
   ```

   b. Check the Solstice Backup boot-time startup file to see whether nsrexecd is being run with the -s option. If the -s option exists, remove all occurrences of -s servername in the file.

   c. Edit or create the /nsr.NetWorker.local/res/servers file:

   – Add the set of Solstice Backup servers, one per line, that you want to grant access to this client.

   – Add an entry for the Solstice Backup logical hostname first. Then add entries for each physical host that can run the Solstice Backup resource group. For example:

   clus_vir1

   clus_phys1

   clus_phys2

d. Restart the Solstice Backup processes:

```
# /etc/init.d/networker start
```

3. On one node in the cluster, start the Solstice Backup service using the cluster management software. For example:

```
# scswitch -Z -g resourcegroup_name
```

4. If required, grant access to the highly available Solstice Backup server on clients outside of the cluster:

   On each client that is outside of the cluster:

   a. Shut down the Solstice Backup processes and verify that all Solstice Backup daemons have stopped:

```
# /usr/sbin/nsr/nsr_shutdown
```

   b. Edit or create the /nsr/res/servers file:

      – Add the set of Solstice Backup servers, one per line, that you want to grant access to this client.

      – Add an entry for the Solstice Backup logical hostname first. Then add entries for each physical host that can run the Solstice Backup resource group. For example:

      clus_vir1

      clus_phys1

      clus_phys2

   c. Restart the Solstice Backup processes:

```
# /etc/init.d/networker start
```

Task 5: Configure the Solstice Backup Server

To configure the Solstice Backup server:

1. Log in as root on the cluster node that is currently running the Solstice Backup resource group.

2. Start the Solstice Backup Administrator program:

```
# nwadmin -s clus_vir1
```

3. Select Server Setup from the main menu.

   a. In the Administrator attribute, add entries for any cluster nodes that are not already listed. For example:

   root@hostname

   b. Click Add and select Apply.

Task 6: Configure Clients under the Solstice Backup Server

When you run the `/usr/sbin/nsr/networker.cluster` script, it creates a symbolic link, `/nsr`, that points to a local disk. It also creates a second link named `nsr.NetWorker.local` that points to the local Solstice Backup directory. For example, if the local Solstice Backup directory was created in `/var/nsr`, each client member will have the following links:

■ `/nsr->/nsr.NetWorker.local`

■ `/nsr.NetWorker.local->/var/nsr`

To configure each client under the Solstice Backup server:

1. From any node in the cluster, start the Solstice Backup Administrator program:

```
# nwadmin -s clus_vir1
```

2. (Optional) Define save groups as required:

   a. Select Customize>Groups from the main menu. The Groups window appears.

   b. Select View>Details from the Groups window.

   c. Click Create to create a new save group.

> ⚠️ **Caution –** Save groups must have the Autorestart attribute enabled and the Manual Restart option disabled in order for their save sets to restart after a virtual client or Solstice Backup server failover.

    d. For each save group, click Enabled in the Autorestart field.

    e. For each save group, make sure that the Manual Restart option is unchecked in the Options check box.

    f. Define the remaining attributes and click Apply.

3. Make each physical client within the cluster a Solstice Backup client of the highly available Solstice Backup server.

   For each physical client in the cluster:

    a. Using the Solstice Backup Administrator program, select Client>Clients Setup.

    b. If you are performing these steps after updating the Solstice Backup server, click on the client name. Otherwise, click Create and enter the name of the physical client for the Name attribute.

    c. In the Remote Access field, add entries for each physical client within the cluster. For example: *root@clus_phys1.*

    d. For the Group attribute, select a group.

    e. Define the remaining attributes in the Client window then click Add, and select Apply.

4. Make each virtual client within the cluster a Solstice Backup client of the highly available Solstice Backup server.

   For each virtual client in the cluster:

    a. Using the Solstice Backup Administrator program, select Client>Clients Setup.

    b. If you are performing these steps after updating the Solstice Backup server, click on the client name. Otherwise, click Create and enter the name of the virtual client for the Name attribute.

    c. In the Remote Access field, add entries for each physical client within the cluster. For example: *root@clus_phys1.*

    d. For the Group attribute, select a group.

    e. Define the remaining attributes in the Client window, click Add, and select Apply.

   The first time the Solstice Backup application runs, it creates the Client resource for the highly available Solstice Backup server.

Task 7: Create Instances of the Solstice Backup Client Resource Type

An instance of the Solstice Backup client resource type must be created for each virtual client that accesses data on globally mounted file systems. Virtual clients in Sun Cluster 3.0 are either logical hostnames or shared addresses.

Caution – All globally mounted file systems (except the /global/.devices/... file systems) must be owned by a logical host and defined in a Solstice Backup client resource type. If the file systems are not properly configured, multiple copies will be backed up for each cluster node.

To back up data for a virtual client, perform the following steps from one node in the cluster:

1. Create an instance of the Solstice Backup client resource as part of the existing resource group that contains the logical host or shared address. For example:

```
# scrgadm -a -j resource_name -g resourcegroup_name -t LGTO.clnt \
-x clientname=virtual_hostname -x owned_paths=pathname_1, pathname_2[,...]
```

The *virtual_hostname* variable is either a logical hostname or shared address. For more information, see the following examples:

- "Example: A Highly Available Informix Database Server" on page 107
- "Example: A Scalable Apache Web Server" on page 107

2. Run a test probe to verify that the client and group resources have been properly configured.

On the node on which the Solstice Backup server resides, enter the following:

```
# savegrp -pv -c client_name savegroup_name
```

If the test probe does not display the scheduled backups and index you want, see "Tracking Scheduled Saves" on page 130.

Example: A Highly Available Informix Database Server

In the following example, the Informix database server is configured to use the DNS registered hostname, *informix_lhrs*. There is an existing failover resource group named *informix_rg* that contains a:

- `SUNW.informix` resource named *informix_res*
- `SUNW.LogicalHostname` resource named *informix_lhrs*

This `SUNW.informix` database server can access data on a global file system under `/global/informix/config` and `/global/informix/db`.

To add a Solstice Backup virtual client to the existing resource group *informix_rg*, enter the following command:

```
# scrgadm -a -j informix_clntrs -g informix_rg -t LGTO.clnt \
-x clientname=informix_lhrs \
-x owned_paths=/global/informix/config,/global/informix/db
```

Example: A Scalable Apache Web Server

In the following example, an Apache web server is configured to use the DNS registered hostname, *apache_sars*. There is an existing scalable resource group named *apache_rg* that contains a:

- `SUNW.apache` resource named *apache_res*
- `SUNW.SharedAddress` resource named *apache_sars*

This Apache web server accesses data on global file system under `/global/web/config` and `/global/web/data`.

To add a Solstice Backup virtual client to the existing resource group *apache_rg*, enter the following command:

```
# scrgadm -a -j apache_clntrs -g apache_rg -t LGTO.clnt \
-x clientname=apache_sars \
-x owned_paths=/global/web/config,/global/web/data
```

Task 8: Register and License the Highly Available Solstice Backup Server

> **Note –** Network Edition does *not* support backup and recover for a highly available Solstice Backup server. For a highly available Solstice Backup server, a Power Edition license is required. Power edition includes two Cluster Client licenses of the same platform type; adding any additional cluster clients will require additional Cluster Client Connection licenses.

To register and license the highly available Solstice Backup server:

1. Order licenses from Sun for the type of Solstice Backup server and client configuration you are using. For more information about cluster license requirements, see "Solstice Backup Cluster Licensing" on page 142

   When you receive the purchased enabler certificate in the mail, proceed to Step 2.

2. Log on to the system that is currently running the Solstice Backup resource group.

3. Perform the following operations to enter the purchased enabler codes:

   a. If you have been evaluating the Solstice Backup highly available software and have now purchased the software:

      i. Upgrade from your base evaluation enabler to your purchased base enabler by entering the following command:

   ```
   # nsrcap -v -u base_enabler_code
   ```

      ii. (Optional) Delete any remaining evaluation enablers using the Solstice Backup Administrator program:

   ```
   # nwadmin -s clus_vir1
   ```

      This will eliminate warning messages due to expiring evaluation enablers.

   b. For each enabler code, enter the following command:

   ```
   # nsrcap -v -c enabler_code
   ```

Task 9: Create a Composite Hostid

To create a composite hostid for the Highly Available Solstice Backup Server:

Caution – A built-in temporary enabler code is in effect for 30 days after installation. If you do *not* enable the software within the 30 day period using either an evaluation or purchased enabler code, you *cannot* back up data.

1. Make sure that you have already defined the highly available Solstice Backup server within the cluster.

2. On each node that can run the highly available Solstice Backup server, run the following command and capture the output:

```
# hostid
```

This prints the numeric identifier of the current host.

3. Log on to the system that is currently running the Solstice Backup resource group.

   a. Create a file named /nsr/res/hostids that contains the hostids of all the cluster nodes that you captured in the previous step. This file points to the global area. Use the following syntax: *hostid1:hostid2:hostid3:...*

   For example: *12345678:87654321*

   b. Restart the server by taking the highly available Solstice Backup server offline and then putting it back online. For example:

```
# scswitch -F -g resourcegroup_name
# scswitch -Z -g resourcegroup_name
```

4. On one node in the cluster, start the Solstice Backup Administrator program using the logical hostname for the highly available Solstice Backup server:

```
# nwadmin -s clus_vir1
```

   a. Select Server>Registration from the main menu.

   b. Note the new hostid that is displayed in the Registration window. This is your composite hostid.

   c. Click Ok to close the Registration window.

Task 10: Authorize the Software

Caution – Do *not* change the logical hostname for the highly available Solstice Backup server. If you change the logical host configuration for the highly-available Solstice Backup server after you update the software, you must permanently license and authorize the highly available Solstice Backup server.

To enter the permanent authorization code for Solstice Backup servers and clients:

1. On one node in the cluster, start the Solstice Backup Administrator program using the logical hostname for the highly available Solstice Backup server:

```
# nwadmin -s clus_vir1
```

2. Select Server>Registration>View>Tabular from the main menu.

   The registration window lists all enabled modules by name, serial number, composite hostid, the expiration date, and the enabler's authorization key.

   – If an enabler *is* authorized, no expiration date is shown.

   – If an enabler is *not* authorized, an authorization key is *not* shown.

3. Select Server>Server Registration from the main menu and then enter the following required information:

   - Name
   - Company
   - Address
   - Phone number
   - Email
   - Date purchased

4. With the registration information and the composite hostid available, contact Sun.

   You can either:

   - Call 1-800-usa4sun
   - Send email to license@sun.com.

   To send email, copy and paste the registration window contents (enabler codes, composite hostid, and expiration dates) along with your name, company name, address, phone number, email address, and date of purchase into an email message. Include a brief note stating that you are requesting authorization keys.

   The authorizations keys are returned within two days by the same method you use to request them.

> ⚠ **Caution –** If you do *not* enter an authorization key within the 45 day period after entering the purchased enabler code, you *cannot* back up data. However, you can recover data.

5. When the authorization keys arrive, perform the following steps:

   a. On one node in the cluster, start the Solstice Backup Administrator program using the logical hostname for the highly available Solstice Backup server:

   ```
   # nwadmin -s clus_vir1
   ```

   b. Select Server>Registration from the main menu.

   c. Enter each authorization key into the Authorization Key field for the corresponding enabler and then click Apply.

▼ To Install Only the Solstice Backup Client Software in a Cluster

In this configuration, the Solstice Backup server is running on a node that is not a member of the cluster.

To install Solstice Backup cluster client, read and follow these procedures:

> ⚠ **Caution –** Make sure that you install the Solstice Backup client software on each node in the cluster.

Task 1: Install the Solstice Backup Cluster Client Software

Note – Before you begin, make sure that the Solaris 8 operating environment and Sun Cluster 3.0 software is already installed on all cluster nodes and that those nodes boot up in the cluster mode.

Also, make sure that your PATH environment variable includes /usr/sbin and /usr/cluster/bin.

To install Solstice Backup software on the computer that you are designating as the Solstice Backup client:

1. Access the Solstice Backup software from the distribution media. For detailed instructions, refer to the *Solstice Backup 6.1 Installation Guide*.

2. Enter the pkgadd command at the system prompt. For example:

```
# /usr/bin/pkgadd -d /cdrom/cdrom0/solaris/sparc
```

The following Solstice Backup software packages are listed:

1. SUNWsbuc (client software package)
2. SUNWsbum (optional man pages)
3. SUNWsbun (storage node software package)
4. SUNWsbus1 (server software package)
5. SUNWsbus2 (device drivers package)

Caution – Do *not* press the [Enter] key for the default response All. Accepting the default installs the server and the device driver packages.

---

3. At the prompt, type **1** to install the client package (SUNWsbuc). For the client installation, only the client package is required.

4. (Optional) At the prompt, type **2** to install the man page package (SUNWsbum).

5. Depending on which version of the Solstice Backup software you are installing, the installation steps differ:

- If you are installing Solstice Backup release 6.1 L10N software:

  a. When prompted to start the Solstice Backup daemons, type **yes**.

  b. When the prompt appears after the Solstice Backup daemons have started, type **q** to exit.

- If you are installing the Solstice Backup release 6.1 with patch 111962-xx:

  a. When prompted to start the Solstice Backup daemons, type **no**.

  b. When you have added all the applicable packages, and the prompt appears type **q** to exit.

  c. Download and install the patch 111962-xx from the Sun web site. For detailed information see, "Downloading and Installing Patch 111962-xx:" on page 113.

Downloading and Installing Patch 111962-xx:

To download and install patch 111962-xx from the Sun web site:

1. In a web browser, enter the following address:

```
http://www.sun.com/service/
```

2. Click the Patches link at the *Online Service Center* section and navigate to the *Sunsolve Online* page.

3. Click the Patchfinder link at the *Sunsolve Contents* section and navigate to the *Sunsolve Online Patches* page.

4. Enter the patch ID 111962 in the text field entry box, click Find Patch, and navigate to the *README* page for the patch.

Note – Patchfinder can find a patch with only its patch id. The revision number -xx is not required.

---

5. Click either the HTTP or FTP link at the top of the page to download the patch. The downloaded patch is in compressed tar format.

6. Uncompress and untar the patch:

```
# cd download_path
# /bin/uncompress 111962-xx.tar.Z
# /bin/tar xvf 111962-xx.tar
```

7. Enter the patchadd command at the system prompt. For example:

```
# cd download_path
# /usr/sbin/patchadd 111962-xx
```

8. Verify that the Solstice Backup daemons are running. For example:

```
# ps -ef | grep nsr
```

9. If the daemons are *not* running, use the following command to start them.

```
# /etc/init.d/networker start
```

Task 2: Configure Solstice Backup Client Software as Highly Available

To define and configure a Solstice Backup client as highly available:

1. Log in as root on each node where you installed the Solstice Backup client software.

2. Make sure that the /etc/hosts file on each cluster node contains the name of the logical host. You can also publish the logical hostname in the Domain Name System (DNS) or Network Information Service (NIS).

3. For each node in the cluster:

   a. Run the cluster configuration script, /usr/sbin/nsr/networker.cluster.

     This script defines the LGTO.clnt resource type that the Solstice Backup software requires.

   b. In response to the prompts, enter the following information:

```
Enter directory where local NetWorker database is installed
[/nsr]?
```

     – Enter the location of the local Solstice Backup database directory provided during the installation procedure. For example: */space/nsr*

```
Do you wish to configure for both NetWorker server and client?
Yes or No [Yes]?
```

     – Enter **no**. This configures only the client software.

Notes:

- The examples used in this procedure are in reference to the examples found in "Required Information: Highly Available Solstice Backup Server" on page 85.
- If you make mistakes during this configuration task, you can undo any changes to the configuration by running the networker.cluster -r option and then running the networker.cluster script again.

Task 3: Create Instances of the Client Resource Type

An instance of the Solstice Backup client resource type must be created for each virtual client that accesses data on the globally mounted file systems. Virtual clients in Sun Cluster 3.0 are either logical hostnames or shared addresses.

Caution – All globally mounted file systems (except the /global/.devices/... file systems) must be owned by a logical host and defined in a Solstice Backup client resource type. If the file systems are not properly configured, multiple copies will be backed up for each cluster node.

To back up data for a virtual client:

1. Create a resource group:

```
# scrgadm -a -g resourcegroup_name
```

2. Add the logical host:

```
# scrgadm -a -L -g resourcegroup_name -l clus_vir1
```

3. Create the client resource:

```
# scrgadm -a -j resource_name -g resourcegroup_name -t LGTO.clnt \
-x clientname=virtual_hostname -x owned_paths=pathname_1, pathname_2[,...]
```

The *virtual_hostname* variable is either a logical hostname or shared address. For more information, see the following examples:

- "Example: A Highly Available Informix Database Server" on page 116
- "Example: A Scalable Apache Web Server" on page 117

Example: A Highly Available Informix Database Server

In the following example, the Informix database server is configured to use the DNS registered hostname, *informix_lhrs*. There is an existing failover resource group named *informix_rg* that contains a:

- SUNW.informix resource named *informix_res*
- SUNW.LogicalHostname resource named *informix_lhrs*

This SUNW.informix database server can access data on a global file system under /global/informix/config and /global/informix/db.

To add a Solstice Backup virtual client to the existing resource group *informix_rg*, enter the following command:

```
# scrgadm -a -j informix_clntrs -g informix_rg -t LGTO.clnt \
-x clientname=informix_lhrs \
-x owned_paths=/global/informix/config,/global/informix/db
```

Example: A Scalable Apache Web Server

In the following example, an Apache web server is configured to use the DNS registered hostname, *apache_sars*. There is an existing scalable resource group named *apache_rg* that contains a:

- `SUNW.apache` resource named *apache_res*
- `SUNW.SharedAddress` resource named *apache_sars*

This Apache web server accesses data on global file system under `/global/web/config` and `/global/web/data`.

To add a Solstice Backup virtual client to the existing resource group *apache_rg*, enter the following command:

```
# scrgadm -a -j apache_clntrs -g apache_rg -t LGTO.clnt \
-x clientname=apache_sars \
-x owned_paths=/global/web/config,/global/web/data
```

Task 4: Define the List of Trusted Solstice Backup Servers

Before a Solstice Backup server can back up a client, the client must grant the server access. Granting access is controlled by the `servers` file. For a node running the Solstice Backup resource group, this file is located in the globally mounted file system. Otherwise, this files is located on a local disk. If the `servers` file does *not* exist, any Solstice Backup server can back up this client. If the `servers` file *does* exist, only the specific servers listed in the file, can back up the client.

To define the list of trusted Solstice Backup servers, perform the following steps on each node in the cluster:

1. Shut down the Solstice Backup processes and verify that all Solstice Backup daemons have stopped:

```
# /usr/sbin/nsr/nsr_shutdown
```

2. Edit or create the `/nsr/res/servers` file and add the set of Solstice Backup servers, one per line, that you want to grant access to this client.

3. Check the Solstice Backup boot-time startup file to see whether `nsrexecd` is being run with the `-s` option. If the `-s` option exists, remove all occurrences of `-s servername` in the file.

4. Restart the Solstice Backup processes:

```
# /etc/init.d/networker start
```

Task 5: Configure Clients under the Solstice Backup Server

To configure each client under the Solstice Backup server:

1. From any node in the cluster, start the Solstice Backup Administrator program:

```
# nwadmin -s server_name
```

2. Make each physical client within the cluster a Solstice Backup client of the Solstice Backup server.

   For each physical client in the cluster:

   a. In the Solstice Backup Administrator program, select Client>Clients Setup.

   b. If you are performing these steps after updating the Solstice Backup server, click on the client name. Otherwise, click Create and enter the name of the physical client for the Name attribute.

   c. In the Remote Access attribute, add entries for each physical client within the cluster. For example: *root@clus_phys1*

   d. For the Group attribute, select a group.

   e. Define the remaining attributes in the Client window, click Add, and select Apply.

3. Make each virtual client within the cluster a client of the Solstice Backup server.

   For each virtual client in the cluster:

   a. In the Solstice Backup Administrator program, select Client>Clients Setup.

   b. If you are performing these steps after updating the Solstice Backup server, click on the client name. Otherwise, click Create and enter the name of the virtual client for the Name attribute.

   c. In the Remote Access attribute, add entries for each physical client within the cluster. For example: *root@clus_phys1*.

   d. For the Group attribute, select a group.

   e. Define the remaining attributes in the Client window, click Add, and select Apply.

4. Schedule backups using the Solstice Backup application. For detailed instructions, refer to the *Solstice Backup 6.1 Administrator's Guide.*

Configuring an External Client to a Highly Available Solstice Backup Server

Use the following procedure to configure and back up an external Solstice Backup client to a highly available Solstice Backup server:

Before a Solstice Backup server can back up a client, the client must grant the server access. Granting access is controlled by the `/nsr/res/servers` file. If this file does *not* exist, any Solstice Backup server can back up this client. If this file *does* exist, only the specific servers listed in the file can back up the client.

On each Solstice Backup client that is outside of the cluster:

1. Shut down the Solstice Backup processes and verify that all Solstice Backup daemons have stopped:

```
# /usr/sbin/nsr/nsr_shutdown
```

2. Edit or create the /nsr/res/servers file:

    a. Add the set of Solstice Backup servers, one per line, that you want to grant access to this client.

    b. Add an entry for the Solstice Backup logical hostname first. Then add entries for each physical host that can run the Solstice Backup resource group.

    For example:

    clus_vir1

    clus_phys1

    clus_phys2

Defining Ownership of a Raw Partition for Virtual Clients

To back up the raw partitions of a Solstice Backup virtual client, you must make sure that the raw device path appears in the owned paths field of the LGTO.clnt resource.

Example: A Highly Available Web Server

In the following example, the highly available web server uses the logical hostname *web_server*. The web server daemons are placed under the control of the resource group *apache_rg*. The daemons access files that are kept in two separate globally mounted file systems: /global/web/config and /global/web/data. The web server also accesses the raw partition /dev/md/hunt/dsk/d30.

To create the Solstice Backup client resource named *apache_nw,* enter the following command:

```
# scrgadm -a -j apache_nw -g apache_rg -t LGTO.clnt \
-x clientname=web_server \
-x owned_paths=\
/global/web/config,/global/web/data,/dev/md/hunt/dsk/d30
```

For more information on backing up raw partitions, refer to the rawasm command as described in the uasm(1m) man page.

Uninstalling the Solstice Backup Software

To uninstall the Solstice Backup software from a cluster environment:

1. On one node in the cluster, disable and remove all instances of the LGTO.serv and LGTO.clnt resource types in the cluster. Refer to the scswitch(1m) and scrgadm(1m) man pages.

2. On each node in the cluster:

   a. Log in as root on the system where you are removing the software.

   b. Undo all changes to the configuration by running the option, networker.cluster -r.

   c. Enter the following command to shut down the Solstice Backup daemons:

   ```
   # /usr/sbin/nsr/nsr_shutdown
   ```

   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.

Caution – The Solstice Backup software packages depend on each other. To remove them all, you must remove them in the following order: SUNWsbus1, SUNWsbun, SUNWsbus2, SUNWsbuc. The man pages SUNWsbum and document files have no dependencies—you can remove them at any time.

d. Enter the following command to remove *all* the packages:

```
# pkgrm SUNWsbus1 SUNWsbun SUNWsbus2 SUNWsbuc SUNWsbum
```

Solstice Backup Cluster Script

This example, generated when you run the **networker.cluster** script, uses the following configuration:

- Logical host: *clus_vir1*
- Mount point: */global/nw*

```
# networker.cluster


--------------------------------------------------------------

Sun Cluster 3.0 extends the Solaris operating environment into a
cluster operating system. A cluster is a collection of loosely
coupled computing nodes that provide a single client view of
network services or applications.

To complete this procedure, one must have a logical host defined
on which the NetWorker server can be entered


--------------------------------------------------------------

Do you wish to continue? [Yes]? yes

Shutting down NetWorker services...

Enter directory where local NetWorker database is installed
[/nsr]? /space/nsr
```

```
Restarting syslog daemon...

Do you wish to configure for both NetWorker server and client?

Yes or No [Yes]? yes

Do you wish to add now the site-specific values for:

NSR_SHARED_DISK_DIR and NSR_SERVICE_ID

Yes or No [Yes]? yes

Enter the Logical Hostname to be used for NetWorker? clus_vir1

The shared directory of the NetWorker server (the nsr directory)
must be on a globally mounted file system.

In what path will this nsr directory be created/located?
/global/nw

Creating /usr/sbin/nsr/lcmap file ...

-rwxrwxr-x   1 root         1707 Aug 20 12:44 /usr/sbin/nsr/lcmap

Finished creating /usr/sbin/nsr/lcmap file.

resource type is not registered

Defining LGTO.clnt resource type with RGM.

resource type is not registered

Defining LGTO.serv resource type with RGM.

NetWorker has been successfully cluster-configured.
```

Solstice Backup for Sun Cluster 3.0 Administration

This section discusses Solstice Backup support for backing up and recovering data that resides on a cluster.

The main sections in this chapter are:

- "Backing Up Data in a Sun Cluster 3.0 Environment" on page 126
- "Recovering Data in a Sun Cluster 3.0 Environment" on page 128
- "Tracking Scheduled Saves" on page 130
- "Configuring an External Client to a Highly Available Solstice Backup Server" on page 134
- "Using a Remote Storage Node with a Highly Available Solstice Backup Server" on page 135
- "Using Tape Libraries and Stand-Alone Tape Devices" on page 136
- "Solstice Backup Cluster Licensing" on page 142

Backing Up Data in a Sun Cluster 3.0 Environment

This section discusses the following topics:

-
-
-

▼ To Configure a Scheduled Save from a Physical Cluster Client

Backing up data from a physical client within a cluster is similar to backing up data on a computer that is not part of a cluster. For detailed instructions on how to back up data, refer to the *Solstice Backup 6.1 Administrator's Guide*.

To prepare for regular backups, install the Solstice Backup client software on all physical clients within the cluster. For installation and configuration instructions, see

▼ To Configure a Scheduled Save from a Virtual Cluster Client

To configure each client under the Solstice Backup server:

1. From any node in the cluster, start the Solstice Backup Administrator program:

```
# nwadmin -s server_name
```

2. (Optional) Define save groups as required.

   a. Select Customize>Groups from the main menu. The Groups window appears.

   b. Select View>Details from the Groups window.

   c. Click Create to create a new save group.

> ⚠️ **Caution –** Save groups must have the Autorestart attribute enabled and the Manual Restart option disabled in order for their save sets to restart after a virtual client or Solstice Backup server failover.

   d. For each save group, click Enabled for the Autorestart attribute.

   e. For each save group, clear the Manual Restart option in the Options check box.

   f. Define the remaining attributes and click Apply.

3. Make each virtual client within the cluster a Solstice Backup client of the highly available Solstice Backup server.

   For each virtual client in the cluster:

   a. In the Solstice Backup Administrator program, select Client>Clients Setup.

   b. Click Create and enter the name of the virtual client for the Name attribute.

   c. For the Remote Access attribute, add entries for each physical client within the cluster. For example:

   root@clus_phys1

   d. For the Group attribute, select a group.

   e. Define the remaining attributes in the Client window and click Apply.

4. Run a test probe to verify that the Client and Group resources have been properly configured.

   On the node on which the Solstice Backup server resides, enter the following command:

```
# savegrp -pv -c virtual_client group_name
```

> ⚠️ **Caution –** If the test probe does not display the scheduled backups and index you want, see "Tracking Scheduled Saves" on page 130.

▼ To Perform a Manual Backup from a Virtual Cluster Client

Backing up data from a virtual cluster client is similar to backing up data from noncluster clients. For example, use one of the following commands:

```
# save saveset
```

or

```
# save -c client saveset
```

Recovering Data in a Sun Cluster 3.0 Environment

This section discusses the following topics:

▼ To Recover Data from a Private Disk Belonging to a Physical Client

Recovering data backed up from a *private disk* belonging to a physical client in a cluster is similar to recovering data on a computer that is not part of a cluster.

For detailed instructions, refer to the *Solstice Backup 6.1 Administrator's Guide*.

▼ To Recover Data from a Shared Disk Belonging to a Virtual Client

To recover data backed up from a *shared disk* belonging to a virtual client:

1. Make sure that you have properly configured the virtual client for remote access from the node that you are currently on.

2. Start the Solstice Backup Administrator program:

```
# nwadmin -s server_name
```

   a. Select Client>Clients Setup from the main menu.

   b. Make sure that the Remote Access attribute of the virtual client for which you want to recover the data contains the name of the physical host you are using to recover the data. For example:

   root@physical_hostname

   For details on how to configure scheduled backups, see "To Configure a Scheduled Save from a Virtual Cluster Client" on page 126.

3. Make sure that the file systems of the virtual client are mounted.

4. Recover data from a shared disk belonging to a virtual client. You can use either the nwrecover GUI or the command line.

   ■ To recover the data using the nwrecover GUI, make sure that you select the virtual client.

   ■ To perform a manual recover, enter the recover command with the -c option at the command prompt. The -c option must be followed by the name of the client you are attempting to recover. For example:

```
# recover -s server_name -c virtual_client
```

   For further information regarding the recover command, refer to the recover man page.

Tracking Scheduled Saves

To verify that the client and group resources have been properly configured, run a
test probe for each client from the node where the Solstice Backup server is running:

```
# savegrp -pv -c client_name group_name
```

If the test probe does not display all the scheduled save sets you want, do *one* of the
following and run a second test probe to verify the configuration:

■ Check the cluster configuration using the cluster management software. If
 necessary, reconfigure the cluster.

■ Make sure that the save sets defined for the client are owned by that client. If
 necessary, redistribute the client save sets to the appropriate client resources.

Caution – Misconfiguration of the cluster resources might cause scheduled save
sets to be dropped from the backup. For more information, see "Path-Ownership
Rules" on page 131.

To override scheduled save rules (not path-ownership rules), run the following
command on each node in the cluster:

```
# touch /usr/sbin/nsr/pathownerignore
```

This command allows any path to be backed up for a client, whether it is owned
by the virtual client or physical node. For detailed information, see "Path-
Ownership Rules" on page 131.

If you used pathownerignore, check that the Solstice Backup scheduled save uses
the correct client index that you want. If it uses the wrong index, you can force save
sets to go to a specific index.

To force save sets to go to a specific index:

1. Start the Solstice Backup Administrator program:

```
# nwadmin -s server_name
```

2. Select Client>Clients Setup from the main menu.

3. For the Backup Command attribute, enter the name of a backup script that contains save -c *client_name*.

For detailed information about the Backup Command attribute, refer to the *Solstice Backup 6.1 Administrator's Guide*.

Path-Ownership Rules

This section discusses the following topics:

About Path-Ownership Rules

In a cluster environment, Solstice Backup software must distinguish between file systems associated with a physical client and file systems managed by a resource group (a virtual client). The criteria used are called the path-ownership rules. These rules determine which client index a save set is written to.

By default, Solstice Backup software might not back up scheduled save sets due to conflicts with path ownership. This prevents a virtual client from writing to multiple indexes. The Solstice Backup software views the client (which owns the file system) as not having matched the client of the current scheduled save set. Writing to multiple indexes might cause recovery issues.

The following conditions cause a file system to be omitted and *not* backed up during a scheduled save:

- If a file system owned by a virtual client is defined in the save set list for a physical client resource.
- If a file system owned by a physical client is defined in the save set list for a virtual client resource.

▼ To Check Path-Ownership Rules

1. Enter the following command on the Solstice Backup server:

```
# savegrp -p -c client_name
```

2. Review which file systems *client_name* owns. This procedure is part of the normal cluster installation setup. For detailed instructions, see "Installing Solstice Backup Software" on page 94.

To test for the existence of misappropriated save sets, run a test probe with the verbose option set. The command output will indicate to which client indexes a save set will be saved. For example:

```
# savegrp -pv -c client_name group_name
```

To ignore Solstice Backup default path-ownership rules, you can create the <nsr_bin>/pathownerignore file. This file causes the Solstice Backup software to back up the file system in question; however, the file system is saved under the index of its correct owner.

If the scheduled save uses the wrong index, you can force save sets to go to a specific index:

1. From any node in the cluster, start the Solstice Backup Administrator program:

```
# nwadmin -s server_name
```

2. Select Clients>Clients Setup from the main menu.

3. For the Backup Command attribute, enter the name of a backup script that contains save –c *client_name*.

   For detailed information about the Backup Command attribute, see the *Solstice Backup 6.1 Administrator's Guide*.

Caution – Creating the <nsr_bin>/pathownerignore file is not recommended, but it might be required under special circumstances.

The <nsr_bin>/pathownerignore file does not override the default path-ownership rules. It causes the path-ownership rules to be ignored when determining if a file system should be backed up during a scheduled save.

▼ To Override Path-Ownership Rules

Overriding the path-ownership rules causes a save set to be saved under a client index different from its default owner.

Manual Save

To override the default path-ownership rule for a manual save, enter the following command:

```
# save -c client_name
```

Scheduled Save

To override the default path-ownership rule for a scheduled save, create a Backup Command attribute that explicitly uses the save -c client_name option.

1. From any node in the cluster, start the Solstice Backup Administrator program:

```
# nwadmin -s server_name
```

2. Select Clients>Clients Setup from the main menu.

3. For the Backup Command attribute, enter the name of a backup script that contains save -c client_name.

   For detailed information about the Backup Command attribute, refer to the *Solstice Backup 6.1 Administrator's Guide*.

For more information, refer to the savefs, save, and savegrp man pages.

Configuring an External Client to a Highly Available Solstice Backup Server

Use the following procedure to configure and back up an external Solstice Backup client to a highly available Solstice Backup server.

Before a Solstice Backup server can back up a client, the client must grant the server access. Granting access is controlled by the `servers` file. For a node currently running the Solstice Backup resource group, this file is located in the globally mounted file system. Otherwise, this file is located on a local disk. If the `servers` file does *not* exist, any Solstice Backup server can back up this client. If the `servers` file *does* exist, only the specific servers listed in the file can back up the client.

On each Solstice Backup client that is outside of the cluster:

1. Shut down the Solstice Backup processes and verify that all Solstice Backup daemons have stopped.

```
# /usr/sbin/nsr/nsr_shutdown
```

2. Edit or create the `/nsr/res/servers` file:

   a. Add the set of Solstice Backup servers, one per line, that you want to grant access to this client.

   b. Add an entry for the Solstice Backup logical hostname first. Then add entries for each physical host that can run the Solstice Backup resource group.

      For example:

      clus_vir1

      clus_phys1

      clus_phys2

3. Restart the Solstice Backup daemons:

```
# /etc/init.d/networker start
```

Using a Remote Storage Node with a Highly Available Solstice Backup Server

This section describes how to access a remote storage node with a highly available Solstice Backup server.

A Solstice Backup storage node can be installed and configured on nodes in the cluster; however, the Solstice Backup storage node is not a highly available application. The following instructions apply in the cases where you are using storage nodes in the cluster and storage nodes that are outside of the cluster.

Caution – Solstice Backup release 6.1 does *not* support highly available remote storage nodes.

To access a remote Solstice Backup storage node from a highly available Solstice Backup server (*rd=physicalhostname:devxxx*):

1. Shut down the highly available Solstice Backup server using the Sun Cluster 3.0 software.

```
# scswitch -F -g resourcegroup_name
```

2. On each node to which the highly available Solstice Backup server can fail over:

   a. Open the Solstice Backup virtual cluster script for editing:
      /usr/sbin/nsr/networker.start

   b. Before the line that starts the nsrd daemons, add the following line:

```
rsh remote_hostname "/usr/sbin/nsr/nsr_shutdown \
-aq && /usr/sbin/nsr/nsrexecd"
```

3. Restart the Solstice Backup server.

Using Tape Libraries and Stand-Alone Tape Devices

This section describes how to use tape libraries and stand-alone tape devices with a highly available Solstice Backup server. Device-sharing infrastructure is defined as hardware, firmware, and software that permit several nodes in a cluster to share access to a device.

Configurations for Tape Libraries and Stand-Alone Tape Devices

If your cluster has a device-sharing infrastructure that supports:

- Both changer and tape devices, use Option 1. For configuration instructions, see "Configuration Option 1" on page 137.

- Changers but does not support tape devices, use Option 2. For configuration instructions, see "Configuration Option 2" on page 139.

You can use Option 3 regardless of your configuration. Option 3 does not require a device-sharing infrastructure. For configuration instructions, see "Configuration Option 3" on page 140.

Caution – Make sure that your device-sharing infrastructure has been certified. For more information, refer to the Legato Compatibility Guides on the Legato web site at `www.legato.com`.

Configuration Option 1

This configuration option offers a seamless backup solution that is 100% hands-free. It allows you to access shared tape devices and shared changer devices with your highly available Solstice Backup server.

Note – Only one tape library is required, although you can add more.
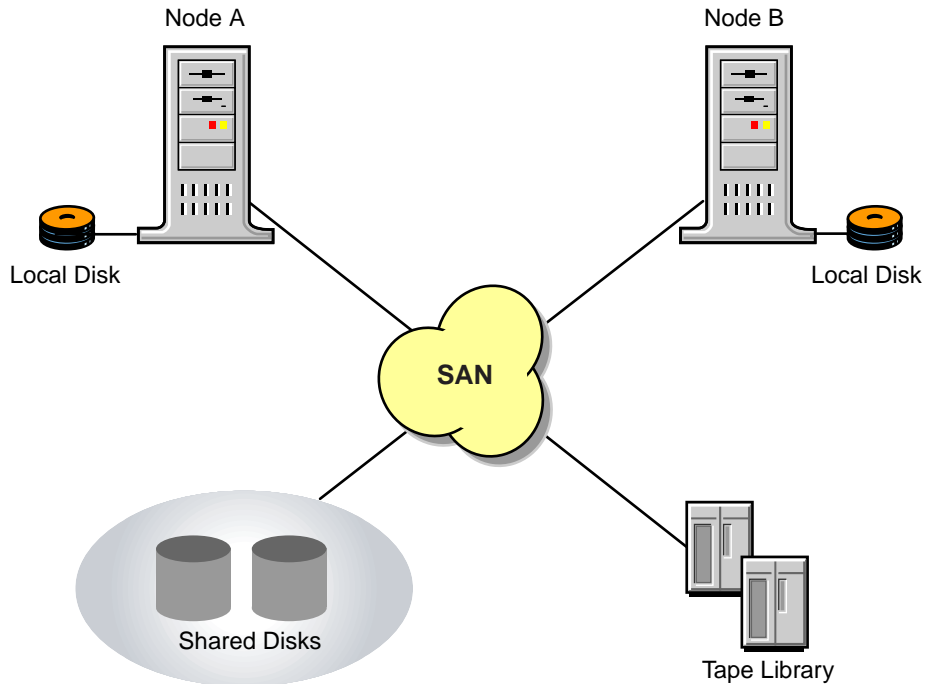
FIGURE 4-1   Configuration Option 1

To configure shared tape devices and shared changer devices to work with your highly available Solstice Backup server:

1. Interface all the tape library devices to all the nodes in the cluster to which the Solstice Backup server can fail over. Use a working SCSI or Fibre Channel switch, and appropriate software as required.

2. Make sure that the shared device infrastructure (SAN bridge, modular data router) supports complete isolation and protection of the path session between the tape library and the node with the Solstice Backup server currently running. The path must be protected from stray bus signals and unauthorized session access from the other nodes.

⚠ **Caution –** If processes on nodes other than the one that is running on the Solstice Backup server can access your tape devices, data corruption might occur. This data corruption might not be detected by the Solstice Backup software.

3. Modify the /usr/sbin/nsr/networker.start script, by adding any required node device-reassignment reservation commands (commands for the bridge) before the line that the starts the nsrd daemon.

4. Test that the device-sharing infrastructure supports path isolation. Use the appropriate operating system commands for your platform against your device, for example, tar, cpio:

   a. From Node A, write a significantly large amount of data to tape. During the write session, power-off and power-on Node B a couple of times.

   b. Verify that the data was properly written by reading back the data from tape and comparing it to the original data.

   If the verification check indicates:

   - Failure, your switching infrastructure does not properly support path isolation. Try Option 2 or Option 3, or contact the device-sharing infrastructure vendor.

   - Success, reverse the node roles:

     – From Node B, write a significantly large amount of data to tape. During the write session, power-off and power-on Node A a couple of times.

     – Verify that the data was properly written by reading back the data from tape and comparing it to the original data.

     – If the verification fails, try Option 2 or Option 3, or contact the device-sharing infrastructure vendor.

5. Use the Solstice Backup jbconfig program to configure the changer and tape devices. For a description of the command, refer to the jbconfig man page.

Configuration Option 2

This option offers a seamless backup solution for accessing nonshared tape devices and shared changer devices from your highly available Solstice Backup server. However, some manual intervention might be required during the data recovery process.

Notes:

- If *N* nodes are used, *N-1* storage node licenses are required. *N* indicates the number of failover nodes for your highly available Solstice Backup server.
- Only one tape library is required, although you can add more.
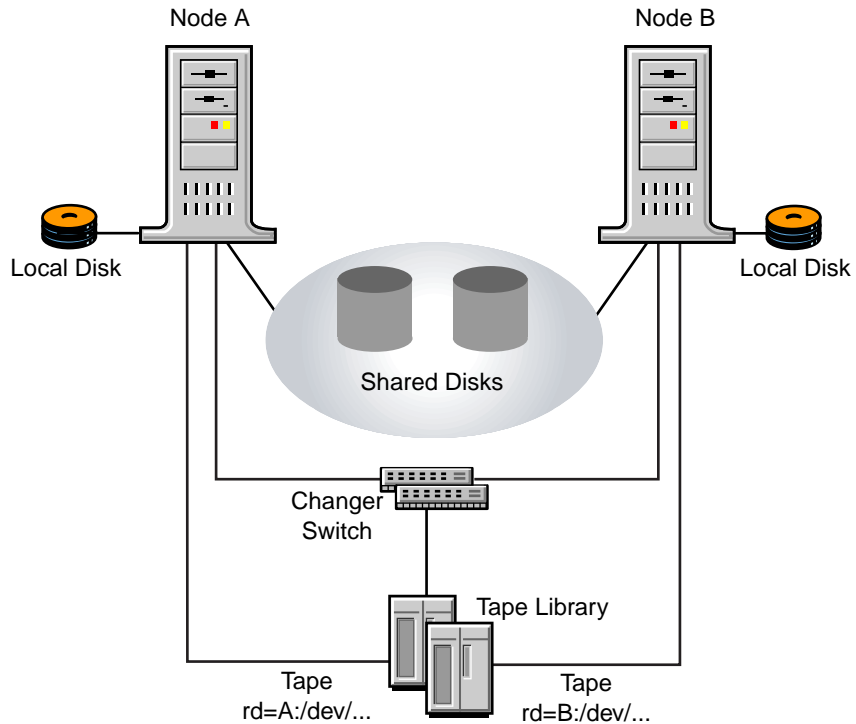
FIGURE 4-2 Configuration Option 2

To configure the devices to work with your highly available Solstice Backup server:

1. Interface the changer device to the Solstice Backup failover nodes. Use a working SCSI or Fibre Channel switch and appropriate software if required.

2. Privately attach each tape device to a separate node within the cluster. For example:

   - Attach tape drive 1 exclusively to Node A.
   - Attach tape drive 2 exclusively to Node B.

3. Use the Solstice Backup jbconfig program to configure the tape devices as remote physical storage nodes. For example:

```
rd="nodeA:/dev/tape#1"
rd="nodeB:/dev/tape#2"
```

Note – When defining each tape device within jbconfig, assign ownership to the node exclusively connected to that tape device. For a detailed description of the jbconfig command, refer to the jbconfig man page.

Configuration Option 3

This option offers a seamless backup solution for accessing nonshared tape devices and nonshared changer devices from your highly available Solstice Backup server. However, some manual intervention might be required during the data recovery process.

Note – *N* indicates the number of failover nodes for your highly available Solstice Backup server.

- If *N* nodes are used, *N-1* storage node licenses are required.
- At least *N* tape libraries are required, one per failover node.
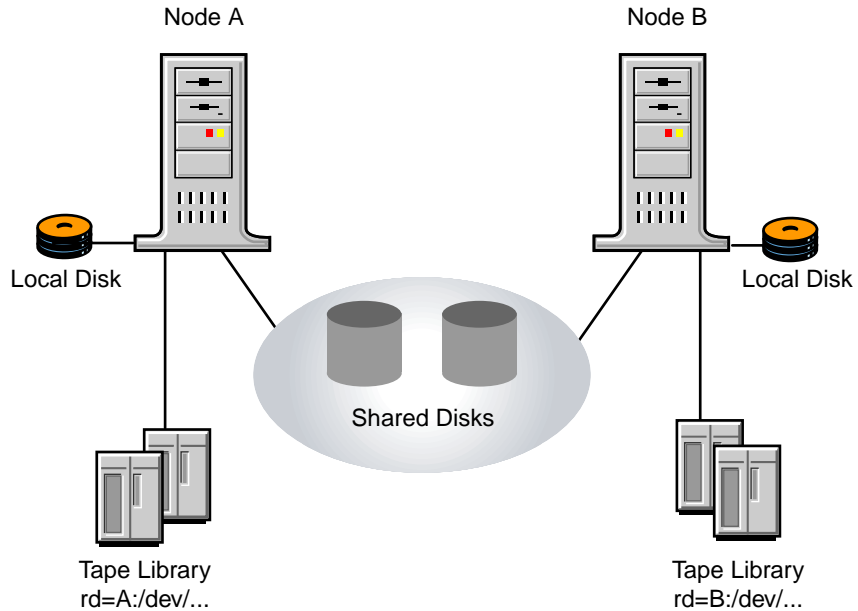
**FIGURE 4-3**   Configuration Option 3

To configure tape changer devices to work with your highly available Solstice
Backup server:

1. Exclusively attach a tape library.

2. Use the following command to configure the nonshared tape and nonshared
   changer devices as remote physical storage node devices. For example:

```
# jbconfig -s server_name
rd="nodeA:/dev/tape#1"
rd="nodeA:/dev/tape#2"
rd="nodeA:/dev/changer"
```

   For a detailed description of the jbconfig command, refer to the jbconfig man
   page.

3. Repeat these steps on each Solstice Backup server failover node.

Solstice Backup Cluster Licensing

This section discusses the following topics:

■ "Enabling the Highly Available Solstice Backup Server" on page 142

■ "Enabling Cluster and Virtual Clients" on page 143

For installation requirements, see "Installing Solstice Backup Software" on page 94.

Enabling the Highly Available Solstice Backup Server

Running the Solstice Backup server as a highly available application on a cluster requires one of the following modes or editions of the Solstice Backup software:

■ Evaluation Mode — unlimited cluster client support during the evaluation period.

■ Power Edition — includes two cluster client licenses of the same platform type; after that, additional Cluster Client Connection licenses are required.

Note – Network Edition does *not* support backup and recover for a highly available Solstice Backup server.

A Cluster Client Connection license is required for each physical node in the cluster on which you intend to run the highly available Solstice Backup server.

In addition, a storage node that is configured locally—using a hostname that matches the physical hostname of the node running the highly available Solstice Backup server—does not require a separate storage node enabler. For example, on a node in a cluster with the physical hostname *Node A*, a remote device defined as *rd=NodeA:/dev/tape1* would not require a separate storage node enabler while the highly available Solstice Backup server was running on *Node A*. An example of this configuration is provided in "Configuration Option 3" on page 140.

Enabling Cluster and Virtual Clients

Solstice Backup client licensing differentiates between stand-alone computers and computers that participate in a cluster. A Solstice Backup server backing up a computer that is a member of a cluster requires a Cluster Client Connection license *instead* of just a Client Connection license.

Cluster Client Connection licenses are distributed on a first-come-first-served basis. Each time a cluster client (physical or virtual) is backed up for the first time, a Cluster Client Connection license is allocated to the physical node where the virtual client is running. Cluster Client Connection licenses are bound to physical nodes. Therefore, once a Cluster Client Connection has been allocated, any virtual clients running on that physical node can be backed up.

Note – Power Edition is only required for a highly available Solstice Backup server. A standalone Solstice Backup server can back up a physical or virtual client of a cluster with a Network Edition license, and the appropriate number of Cluster Client Connection licenses.

Cluster Clients

To back up a cluster client, your Solstice Backup server requires:

- A Client resource defined for that member of the cluster
- A Cluster Client Connection license

Once the Cluster Client Connection license count has been exhausted, you must either:

- Purchase more Cluster Client Connection licenses
- Remove Solstice Backup cluster client resources that have assumed a Cluster Client Connection license

---

Note – If you remove cluster client resources, you *cannot* back up those cluster clients or any virtual clients running on that cluster client.

If a remote cluster client does *not* have the same architecture as the Solstice Backup server, you will need an appropriate ClientPak enabler and Cluster Client Connection license. For example, if the highly available Solstice Backup server is on a Sun Cluster 3.0 system and the cluster client is a member of an Microsoft Cluster Service (MSCS) environment, you will need an NT ClientPak enabler and a Cluster Client Connection license for Windows.

Virtual Clients

To back up virtual clients, the cluster node that is running the virtual client must be licensed within the Solstice Backup server. Cluster Client Connection licenses are bound to the physical node. Only one Cluster Client Connection license is required to back up one or more virtual clients on a single physical node. However, each node within the cluster that can back up virtual clients requires a Cluster Client Connection license.

International (Localized) Versions of Solstice Backup Software

This appendix covers topics pertinent to versions of Solstice Backup software that have been "localized" to languages other than English.

Installing Localized Solstice Backup Software

You must first install the English-language version of Solstice Backup release 6.1 before installing the following non-English versions of:

- French
- Chinese
- Japanese

If you currently run the English-language version of Solstice Backup release 6.1, you can directly install the required localized software.

▼ To Update from Solstice Backup Release 6.0L to Release 6.1

To update to Solstice Backup release 6.1 from release 6.0L:

1. Remove the earlier Solstice Backup release (but *not* the indexes, resource files, and media database).

2. Install the English-language version of Solstice Backup release 6.1. When you have completed that installation, install the appropriate language version of Solstice Backup release 6.1.

3. Enable and register the Solstice Backup software.

For more information, refer to the *Solstice Backup Installation Guide Release 6.1*.

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 can include the forward slash (/), 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 the Solstice Backup software support of locales and code sets:

- "Locale Support" on page 147
- "Code Set Support" on page 147

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

The Solstice Backup software display of non-English file names varies depending on whether the software is localized:

- A localized version of the Solstice Backup software on a localized operating system running a supported locale displays localized filenames correctly.
- An English version of the 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

Solstice Backup releases 6.0L and 6.1 support the following code sets:

- Microsoft Windows NT/Windows 2000:

  Solstice Backup software for Microsoft Windows NT supports the standard Microsoft Windows NT code set, which is 16-bit Unicode. The Japanese version of Microsoft Windows NT and Solstice Backup software uses Shift JIS character encoding.

- Solaris

TABLE A-1 indicates the Solaris Language environment settings for each supported code set.

TABLE A-1   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

Server/Client Compatibility with Localized Solstice Backup Software

The optimal configuration for running the localized Solstice Backup software is the following:

- Localized server (same release as clients)
- Localized clients (same release as server)
- Localized operating system that supports this localized release of Solstice Backup software

The following tables explain:

- Server/client compatibility for localized Solstice Backup software
- Operating system version compatibility with server and client support

For compatibility information regarding your localized version of Solstice Backup software, see one of the following sections:

- TABLE A-2 on page 148
- TABLE A-3 on page 149
- TABLE A-4 on page 149

TABLE A-2   Server/Client Compatibility — French UNIX

Server/Client Compatibility on French Solaris Operating System	
Solaris 8 Operating System — French	
Solstice Backup Server	**Supported Clients**
Solstice Backup Server 6.0L or 6.1	Solstice Backup Client 6.0L or 6.1 French
Solaris 7 Operating System — French	
Solstice Backup Server	**Supported Clients**
Solstice Backup Server 6.0L or 6.1 French	Solstice Backup Client French 6.0L or 6.1 on Solaris 7, 8
Solaris 6 Operating System — French	
Solstice Backup Server	**Supported Clients**
Solstice Backup Server 6.0L or 6.1 French	Solstice Backup Client French on Solaris 6, 7, 8

TABLE A-3   Server/Client Compatibility — Japanese UNIX

Server/Client Compatibility on Solaris Operating System	
Solaris 8 Operating System — Japanese	
Solstice Backup Server	**Supported Clients**
Solstice Backup Server 6.0L or 6.1 Japanese	Solstice Backup Client 6.0L, 6.1, 5.5/5.5.1 Japanese
Solaris 7 Operating System — Japanese	
Solstice Backup Server	**Supported Clients**
Solstice Backup Server 6.0L or 6.1 Japanese	Solstice Backup Client 6.0L, 6.1, 5.5/5.5.1 Japanese
Solstice Backup Server 5.5/5.5.1 Japanese	Solstice Backup Client 5.5/5.5.1 Japanese
Solaris 6 Operating System — Japanese	
Solstice Backup Server	**Supported Clients**
Solstice Backup Server 6.0L or 6.1 Japanese	Solstice Backup Client 6.0L 6.1, 5.5/5.5.1 Japanese
Solstice Backup Server 5.5/5.5.1 Japanese	Solstice Backup Client 5.5/5.5.1 Japanese

TABLE A-4   Server/Client Compatibility — Simplified Chinese

Server/Client Compatibility on Solaris Operating System	
Solaris 8 Operating System — Chinese	
Solstice Backup Server	**Supported Clients**
Solstice Backup Server 6.0L or 6.1 Chinese	Solstice Backup Client 6.0L or 6.1 Chinese
Solaris 7 Operating System — Chinese	
Solstice Backup Server	**Supported Clients**
Solstice Backup Server 6.0L or 6.1 Chinese	Solstice Backup Client 6.0L or 6.1 Chinese
Solaris 6 Operating System — Chinese	
Solstice Backup Server	**Supported Clients**
Solstice Backup Server 6.0L or 6.1 Chinese	Solstice Backup Client 6.0L or 6.1 Chinese

Limitations Specific to Localized (Non-English) Versions of Solstice Backup Software

This section describes limitations in the various localized versions of Solstice Backup software.

Remote Login Might Display Unreadable Messages (Legato Bug ID: LGTpa25309)

In either the Simplified Chinese or French version of Solstice Backup software, during a remote login you might encounter the following problems:

■ When you complete the remote login from computer_A to computer_B, garbled messages appear in the GUI of computer_A.

■ The menus are displayed in English, despite having used a Chinese or French version of Solstice Backup software.

This problem occurs even if the locale settings on both computers match. To work around this situation, complete the following steps:

1. On the computer from which you want to complete the remote login (in this case, computer_A), launch the `nwadmin` GUI.

2. From the Server menu, select Change Server.

3. Enter the name of the server to which you want to log in remotely (for example, computer_B).

Note: With the Japanese version of Solstice Backup software, you can avoid unreadable messages during a remote login by ensuring that the following variable is set as follows:

```
LANG=ja
```

Required Locale Settings in the French Version of Solstice Backup Software (Legato Bug ID: LGTpa25104)

In the French version of Solstice Backup software, ensure that the following locales are set appropriately:

```
LC_MESSAGES=fr
```

LANG may be set to any one of the following:

```
 * fr_FR.ISO8859-1
```

```
 * fr_FR.ISO8859-15
```

```
 * fr
```

If you do *not* set these variables appropriately, the resulting GUI and command line messages might be unreadable.

Localized Software Contains Some English GUIs and Messages (Legato Bug ID: LGTpa32041)

The localized versions of the Solstice Backup software have several components that have *not* been localized. These English language components do *not* affect the functionality of the software.

Solstice Backup for Solaris

The following components of the Solstice Backup software for Solaris have *not* been localized and still appear in English:

- nsrwatch, nsradmin (feature English menu names and help messages)
- jbverify

Unable to load cleaning tape from command line (Legato Bug ID: LGTpa30245)

To successfully load a cleaning tape from the command line with the following command, you must ensure the term cleaning tape is enclosed within quotes:

```
# nsrjb -l -f /dev/nrst1 cleaning tape
```

-R Option with recover Command Not Currently Supported with Solstice Backup 6.1(Legato Bug ID: LGTpa32685)

The -R option with the recover command is not currently supported with any localized version of Solstice Backup 6.1.

Length Limits to File Paths and Directory Names (Legato Bug IDs: LGTpa25330, 25365)

The Solstice Backup 6.x software supports a full file path length limit of 1024 bytes (standard for most operating systems). Consequently, if you are working in a language that uses multi-byte characters (such as, Chinese or Japanese), be sure to limit the file path and directory names to remain within the limits imposed by your operating system and the Solstice Backup software.

Setting LANG Environment Variable after Upgrading to Localized Solstice Backup Software (Legato Bug ID: LGTpa25878)

After upgrading to a localized version of Solstice Backup release 6.0L or later from an earlier, English-only version of the Solstice Backup software, you should initially set the LANG environment variable to C (LANG C). This is important mainly in the Solaris operating system.

Setting the environment variable LANG C allows the recovery of localized files that were backed up before you upgraded to a localized version of Solstice Backup software.

Complete the recoveries from the command line using the `recover` command, rather than through the GUI. If you use the GUI to complete the recovery, you might encounter display problems in the file names, although the actual data that has been recovered is fine.

Non-Localized Solstice Backup Software on Localized Operating Systems Could Have File Display Limitations (Legato Bug ID: LGTpa26051)

An English version of Solstice Backup software on a localized version of an operating system can perform backups and recoveries, but the filename might *not* appear correctly in the Solstice Backup software. A localized version of Solstice Backup software on a localized operating system for the supported locales backs up, recovers, and displays the files correctly.

English Solstice Backup software on an English-language operating system supports backup, recovery, and display of English filenames only.

Note – This issue does *not* occur if you are using an English-language version of Solstice Backup software with a Chinese version of Microsoft Windows.

Restoring nsr/gls Directory on Localized Solstice Backup Software (Legato Bug ID: LGTpa26972)

If you are using a localized version of Solstice Backup software and lose or damage any files under the /nsr/gls directories, you will experience problems with the Solstice Backup displays. The /nsr/gls directory is required to display the localized version of the Solstice Backup software.

To restore the localized version of Solstice Backup software successfully, do one of the following:

- Recover the /nsr/gls directory separately.
- Remove and then reinstall the localized Solstice Backup software (SUNWfrsbu, SUNWjasbu, SUNWzhsbu).

Localized Solstice Backup Packages (Legato Bug ID: LGTpa25543)

When using a localized version of the Solstice Backup software, *do not* relocate any of these packages:

- SUNWfrsbu: French
- SUNWjasbu: Japanese
- SUNWzhsbu: Simplified Chinese

Relocating these packages causes problems with the Solstice Backup software displays.

You can, however, relocate the following packages:

- SUNWsbuc
- SUNWsbum (including localized man pages)
- SUNWsbun
- SUNWsbus1

Man Command Dumps Core on LGTfrman on Solaris 2.6 (Legato Bug ID: LGTpa27278)

When using the localized French version of the Solstice Backup software on Solaris 2.6, the `man` command dumps core if you attempt to view or print the localized French man pages.

To work around this limitation, copy `nroff` from another release of the Solaris operating systems, such as 2.5.1, 2.7, or 2.8. This allows you to view or print the French man pages.

For further information about this Sun bug (bug ID4171302) contact Sun Microsystems. Support contract customers can obtain the patch from their local Solution Centers, which as of the date of publication of this document were listed at:

```
http://www.sun.com/service/contacting/solution.html
```

Support contract customers in North America can get the patch information by calling 1-800-USA-4SUN.

The nsrwatch Program Displays Trailing Characters (Legato Bug ID: LGTpa22291)

The `nsrwatch` program in the Japanese and Chinese versions of Solstice Backup software displays the trailing character "|" in the `nrswatch` display window. The trailing character does *not* affect the function of the program.

Mapping Logical Device Paths to Physical Drives

Before sharing drives using DDS or configuring an autochanger, determine the correct mapping of device paths (created by the operating system) to the physical drives.

To determine the correct mapping of the device paths, you must obtain the SCSI addresses for the drives from the hardware. Then you must obtain the mapping of SCSI addresses to the devices paths for each host on the computer.

With this information, you can ensure that the devices are appropriately mapped to the correct drive.

1. Map logical device paths obtain and record the SCSI address for each tape drive.

   Refer to the hardware documentation that explains how to obtain the SCSI address for a specified drive.

2. Run the `inquire` program on each host being configured to obtain a list of devices attached to the computer, along with the corresponding SCSI addresses.

   The `inquire` program displays device information in a format similar to the following:

```
scsidev@4.1.0:STK       9738      1002|Autochanger (Jukebox),

/dev/rac/c4t1d0

scsidev@4.2.0:STK       9840      1.25|Tape, /dev/rmt/2hnb

scsidev@4.3.0:STK       9840      1.25|Tape, /dev/rmt/1hnb
```

3. Using the SCSI address, map the physical tape drives to a corresponding device (created by the operating system) when you configure DDS and/or run the `jbconfig` program.

4. Run the `jbconfig` program.

5. Indicate that the drive will be shared when prompted by the `jbconfig` program.

6. Enter the device path (created by the operating system) for the host that accesses drive 1.

   In the example "Example: Mapping Drives to Devices" on page 158, the server, `neptune, with the device path /dev/rmt/2hnb is the device sharing drive 1.`

7. Complete all prompts as required for server `neptune`.

8. Enter the host name and the device path of the second host that will access drive 1.

   In the example, the storage node, sn_1, with the device path `/dev/rmt/0cbn is the second host sharing drive 1.`

9. Complete all prompts as required for your computer.

Example: Mapping Drives to Devices

In this example, there is an autochanger with two drives that have the following SCSI IDs:

Drive 1: SCSI ID 2 - STK 9840

Drive 2: SCSI ID 3 - STK 9840

You want two different nodes to share drive 1. Consequently, to share drive 1 between two nodes, you enter—when prompted by the `jbconfig` program—the device path that has the same SCSI address as the drive you want shared.

The points at which to enter the logical device paths are noted in the following sample jbconfig program output. In this sample output, Drive 1 is to be shared by:

- A server, neptune — /dev/rmt/2hnb (device path created by operating system)
- A storage node, sn_1 — /dev/rmt/0cbn (device path created by operating 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@4.1.0: Standard SCSI Jukebox, Vendor <STK>, Product
<9738>

Which one do you want to install? 1

Installing an 'Standard SCSI Jukebox' jukebox.

Name you would like to assign to the jukebox device?

STK_9738

A brief description of the jukebox device?

scsidev@4.1.0: stk 9738

Pathname of the control port for the jukebox device? scsidev@4.1.0

Do you want automated device cleaning support enabled? (yes/no) n

Will media drive 1 be shared by multiple device paths? (yes/no) y
```

```
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: ? neptune
```

Note – At the following prompt, enter the device path for media drive 1 on the host
neptune, previously mapped to a physical drive based on the SCSI ID.

```
Enter pathname of device 1 for media drive 1: ?

/dev/rmt/2hnb

using '/dev/rmt/2hnb' 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/2hnb.

        16) 8mm Mammoth-2                        37) tz87
        17) 9840                                 38) tz88
        18) 9940                                 39) tz89


Choice? 17

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 also accesses this drive.

```
Enter hostname that owns device path 2 for media drive 1: ? sn_1
```

Note – The **inquire** command has been run on storage node sn_1 and cross referenced to the known physical SCSI ID hardware map to correctly determine the device path for media drive 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.

        16) 8mm Mammoth-2                      37) tz87
        17) 9840                               38) tz88
        18) 9940                               39) tz89

Choice? 17

Any more device paths for media drive 1? (yes/no) n

Will media drive 2 be shared by multiple device paths? (

yes/no) n

Enter hostname that owns media drive 2: ? neptune
```

Note – At the following prompt, enter the device path for media drive 2 on the host neptune, previously mapped to a physical drive based on the SCSI ID.

```
Enter pathname of device 1 for media drive 2:

/dev/rmt/1hnb

using '/dev/rmt/1hnb' 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/1hnb.

        16) 8mm Mammoth-2                        37) tz87
        17) 9840                                 38) tz88
        18) 9940                                 39) tz89

Choice? 17

Your jukebox has a bar code 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
```

Solstice Backup 6.1 Updating Strategies

This appendix contains the following sections:

- "Migration Path" on page 163
- "Installing Solstice Backup Software" on page 164
- "Additional Information" on page 165

Migration Path

Solstice Backup 5.1 and 5.1.1 releases have different server software product families than those available in the Solstice Backup 5.5 and later release. See the "Migration Path" table for appropriate migration path(s):

TABLE C-1   Migration Path *(1 of 2)*

From Release	From Server Software Product	To 6.1 Server Software Product	License
5.1, 5.1.1	Backup Server Edition	WorkGroup Edition	Update/7 enabler in this media kit + Order additional 6.1 WorkGroup Edition license (LSBIS-610-2106)
	Backup Network Edition	Network Edition	Update/7 enabler in this media kit
	Backup Power Edition	Power Edition	Update/7 enabler in this media kit

TABLE C-1   Migration Path *(2 of 2)*

From Release	From Server Software Product	To 6.1 Server Software Product	License
5.5 (5.5.1)	WorkGroup Edition	WorkGroup Edition	Update/7 enabler in this media kit
	Network Edition	Network Edition	Update/7 enabler in this media kit
	Power Edition	Power Edition	Update/7 enabler in this media kit
6.0	WorkGroup Edition	WorkGroup Edition	Do not need update enabler
	Network Edition	Network Edition	Do not need update enabler
	Power Edition	Power Edition	Do not need update enabler

Installing Solstice Backup Software

For information on installing the Solstice Backup software refer to "Installing Solstice Backup in a Sun Cluster 3.0 Environment" on page 79.

For further information see the *Solstice Backup 6.1 Installation Guide*. This document is available on the Solstice Backup 6.1 CD-ROM and in the Solstice Backup 6.1 documentation kit.

Additional Information

The following section contains additional information when updating to Solstice Backup release 6.1.

Server Jukebox Resource Variables

After you update from the Solstice Backup 5.1/5.1.1 to 6.1 software, the server's jukebox resource (configured for use by the storage node) contains a new set of attributes. For a description of these new attributes and values, refer to "New Attributes Visible in the Jukebox Resource" in *Release Supplement of the Solstice Backup 5.5 Installation Guide and Release Notes*.

This document is available for download from the following public web site:

> *http://www.sun.com/products-n-solutions/hardware/docs/Software/Storage_Software/Solstice_Backup/index.html*

Client File Index and Media Database Format

When you update from the Solstice Backup release 5.x to 6.x software, the client file indexes and media database must be converted to a new format. For details on the conversion procedure, refer to "Installing Solstice Backup in a Sun Cluster 3.0 Environment" on page 79.

For further information see the *Solstice Backup 6.1 Installation Guide*. This document is available on the Solstice Backup 6.1 CD-ROM and in the Solstice Backup 6.1 documentation kit.

Enterprise Application Module 2.x/3.x for Oracle

If you update from Solstice Backup release 5.5 (5.5.1) to 6.x, and you have installed the Enterprise Application Module (EAM) 2.x for Oracle to back up Oracle clients, you can no longer back up Oracle 8x clients. This is because Solstice Backup 5.5 offered one module, EAM 2.x for Oracle for both Oracle 7.x and 8x support. With Solstice Backup 6.0 or later, two modules provide Oracle support:

■ Enterprise Application Module 2.x for Oracle 7.x

■ Enterprise Application Module 3.x for Oracle 8, 8i

You must choose one of the options below if you want to continue backing up Oracle 8x clients after the update:

■ Keep the Enterprise Application Module 2.x for Oracle 7.x (LSBIS-XXX-2151), and purchase the Enterprise Application Module 3.x for Oracle 8, 8i (LSBIS-610-2154).

or

■ Give up Enterprise Application Module 2.x for Oracle (LSBIS-XXX-2151) in exchange for Enterprise Application Module 3.x for Oracle 8x (LSBIS-610-2154) at no charge.

Cluster Client Connection Licenses

Prior to Solstice Backup 6.0, you were able to back up clients that belong to a cluster by using the Client Connection license. A new license was introduced in Solstice Backup release 6.0 and later for backing up clients that belonged to a cluster. When backing up a client, the Solstice Backup release 6.0 software and later determines whether a client belongs to a cluster. If the client belongs to a cluster, it will check for the new Cluster Client Connection license. The absence of a Cluster Client Connection license results in the following error message:

```
UNIX CLUSTER CLIENT license: Too many clients
```

To back up clustered clients using the Solstice Backup release 6.0 or later, customers must purchase the Cluster Client Connection License. This license cannot be used with Workgroup Edition.

Note – The Solstice Backup Power Edition software comes with a Cluster Client Connection license which allows for two cluster client connections.

The new Cluster Client Connection licenses are:

- LSBIS-610-2114 Cluster Client Connection License for UNIX, Quantity 1
- LSBVS-610-2115 Cluster Client Connection License for UNIX, Quantity 5
- LSBYN-610-2110 Cluster Client Connection License for NT, Quantity 1
- LSBYN-610-2111 Cluster Client Connection License for NT, Quantity 5
- LSBIS-610-7058 Cluster Client Connection License for Linux, Quantity 1
- LSBIN-610-7059 Cluster Client Connection License for Linux, Quantity 5

To continue support while you obtain a permanent Cluster Client Connection enabler, use the following "45 Day Temporary Enabler Codes" for UNIX and Windows.

TABLE C-2   45 Day Temporary Enabler Codes

Cluster Client Connection	Temporary Enabler Code
UNIX	Enabler 1: 5af0dc-1ca682-f4b520 Enabler 2: 5df5df-2199bf-f7b627
Windows	Enabler 1: 65cce7-2991b7-ff522f Enabler 2: 66cde8-2892b6-f8ad2c

Linux ClientPak 45 Day Enabler Codes

Prior to Solstice Backup 6.0, Linux client support was included with the ClientPak for UNIX license. As of Solstice Backup 6.0, Solstice Backup for Linux client software is licensed as a separate platform. The ClientPak for Unix product will no longer include Linux support.

To continue Linux support while you obtain a permanent ClientPak for Linux enabler, use the following 45-day temporary ClientPak for Linux enablers:

- Temp Enabler 1: f67378-9f0226-6818bc
- Temp Enabler 2: d25754-b32e0a-794f98

If you wish to have Linux client support, you will need to purchase the ClientPak for Linux license (LSBIX-610-3321).