



Sun Cluster 3.0 12/01 Release Notes Supplement

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054 U.S.A.
650-960-1300

Part No. 816-3753-10
January 2002, Revision A

Copyright 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, CA 95054, U.S.A. All rights reserved.

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

This product or document is distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Java, Netra, Solaris, Sun StorEdge, iPlanet, Sun Cluster, Sun Plex, AnswerBook2, docs.sun.com, Solstice DiskSuite, Sun Enterprise, Sun Enterprise SyMON, Solaris JumpStart, JumpStart, Sun Management Center, OpenBoot, Sun Fire, SunSolve, SunSwift, the 100% Pure Java logo, the AnswerBook logo, the Netra logo, the Solaris logo and the iPlanet logo are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon architecture developed by Sun Microsystems, Inc.

ORACLE® is a registered trademark of Oracle Corporation. Netscape™ is a trademark or registered trademark of Netscape Communications Corporation in the United States and other countries. The Adobe® logo is a registered trademark of Adobe Systems, Incorporated.

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>).

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, CA 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. a les droits de propriété intellectuels relatants à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuels peuvent inclure un ou plus des brevets américains énumérés à <http://www.sun.com/patents> et un ou les brevets plus supplémentaires ou les applications de brevet en attente dans les Etats - Unis et dans les autres pays.

Ce produit ou document est distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, Java, Netra, Solaris, Sun StorEdge, iPlanet, Sun Cluster, Sun Plex, AnswerBook2, docs.sun.com, Solstice DiskSuite, Sun Enterprise, Sun Enterprise SyMON, Solaris JumpStart, JumpStart, Sun Management Center, OpenBoot, Sun Fire, SunSolve, SunSwift, le logo 100% Pure Java, le logo AnswerBook, le logo Netra, le logo Solaris et le logo iPlanet sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

ORACLE® est une marque déposée registre de Oracle Corporation. Netscape™ est une marque de Netscape Communications Corporation aux Etats-Unis et dans d'autres pays. Le logo Adobe® est une marque déposée de Adobe Systems, Incorporated.

Ce produit inclut le logiciel développé par la base de Apache Software Foundation (<http://www.apache.org/>).

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.



Adobe PostScript

Sun Cluster 3.0 12/01 Release Notes Supplement

This document supplements the standard user documentation, including the *Sun Cluster 3.0 12/01 Release Notes* shipped with the Sun™ Cluster 3.0 product. These “online release notes” provide the most current information on the Sun Cluster 3.0 product. This document includes the following information.

- [“Revision Record” on page 4](#)
- [“New Features” on page 5](#)
- [“Restrictions and Requirements” on page 5](#)
- [“Known Problems” on page 6](#)
- [“Known Documentation Problems” on page 8](#)

Revision Record

The following table lists the information contained in this document and provides the revision date for this information.

TABLE 1-1 Sun Cluster 3.0 12/01 Release Notes Supplement

Revision Date	Information Added
January 2002	<p>“Bug ID 4362925” on page 6, problem during Sun Cluster shutdown when <code>scshutdown</code> tries to unmount a cluster file system that the process <code>nsrmmmd</code> is still referencing.</p> <p>“Bug ID 4626010” on page 7, description of known problem and clustering fix regarding the default hardware configuration of Sun StorEdge™ 3900 Series systems.</p> <p>Known documentation problems with the <i>Sun Cluster 3.0 12/01 Hardware Guide</i>. See “Hardware Guide” on page 8.</p> <p>Procedures to support Sun Cluster HA for SAP as a scalable application. See Appendix B.</p> <p>Procedures to support Sun StorEdge 3910 and Sun StorEdge 3960 systems. See Appendix C.</p> <p>Procedures to support Sun StorEdge 9910 and Sun StorEdge 9960 systems. See Appendix D.</p>
December 2001	<p>New Patch 110651-06 replaces 110651-04 to provide Oracle 9i support on Sun Cluster software. See “Data Service Restrictions and Requirements” on page 5.</p> <p>“Bug ID 4525403” on page 7, required Sun Cluster setting in the <code>/etc/system</code> file overwritten when VxFS is installed after Sun Cluster software.</p> <p>Generic data services (GDS) overview and instructions on creating a service that uses GDS is provided in Appendix A.</p>

New Features

In addition to features documented in *Sun Cluster 3.0 12/01 Release Notes*, this release now includes support for the following features.

Generic Data Service

[Appendix A](#) provides information on the generic data service (GDS) and shows you how to create a service that uses GDS using either the command line interface or the SunPlex Agent Builder.

Restrictions and Requirements

The following restrictions and requirements have been added or updated since the Sun Cluster 3.0 7/01 release.

Data Service Restrictions and Requirements

- **Oracle 9i** – Sun Cluster 3.0 software requires Patch 110651-06 or higher to support Oracle 9i. Patch 110651-06 replaces 110651-04, which was found to break Oracle 8i functionality. Read the special installation instructions that accompany 110651-06 before you install the patch.

Known Problems

In addition to known problems documented in *Sun Cluster 3.0 12/01 Release Notes*, the following known problems affect the operation of the Sun Cluster 3.0 12/01 release.

Bug ID 4362925

Problem Summary:

```
nodeA# scshutdow -g0 -y
scshutdow: Unmount of /dev/md/sc/dsk/d30 failed: Device busy.
scshutdow: Could not unmount all PxFS filesystems.
```

The Networker packages were bundled and installed during the Oracle installation. Therefore, the `nsrmmmd` daemon is running and mounting the `/global/oracle` directory, which prevents the unmount of all cluster file systems.

```
nodeA# umount /global/oracle
umount: global/oracle busy
nodeA# fuser -c /global/oracle
/global/oracle: 335co 317co 302co 273co 272co
nodeA# ps -ef|grep 335
root 335 273 0 17:17:41 ?          0:00 /usr/sbin/nsrmmmd -n 1
root 448 397 0 17:19:37 console 0:00 grep 335
```

This problem occurs during Sun Cluster shutdown when the shutdown tries to unmount a cluster file system that the process `nsrmmmd` is still referencing.

Workaround: Run the `fuser(1M)` command on each node to establish a list of all processes still using the cluster file systems that cannot be unmounted. Check that no Resource Group Manager resources have been restarted since the failed `scshutdow(1M)` command was first run. Kill all these processes with the `kill -9` command. This kill list should not include any processes under the control of the Resource Group Manager. After all such processes have terminated, rerun the `scshutdow` command, and the shutdown should run to successful completion.

Bug ID 4525403

Problem Summary: VxFS installation adds the following line to the `/etc/system` file:

```
* vxfs_START -- do not remove the following lines:
*
* VxFS requires a stack size greater than the default 8K.
* The following values allow the kernel stack size
* for nfs threads to be increased to 16K.
*
set rpcmod:svc_default_stksize=0x4000
* vxfs_END
```

If you install VxFS after Sun Cluster software has been installed, the VxFS installation overrides the following required Sun Cluster 3.0 12/01 setting in the `/etc/system` file:

```
* Start of lines added by SUNWscr
set rpcmod:svc_default_stksize=0x6000
```

Workaround: Change the setting value from `0x4000` to `0x6000`.

```
set rpcmod:svc_default_stksize=0x6000
```

Bug ID 4626010

Problem Summary: In a Sun StorEdge 3900 Series system, the preconfigured, default hard zones that are used on the Sun StorEdge Network FC Switch-8 and Switch-16 switches are incompatible for use with host-based mirroring in a cluster. The hard zones prevent the cluster nodes from seeing all attached storage because there are only two ports per zone.

Workaround: Remove the default hard zones from the Sun StorEdge Network FC Switch-8 and Switch-16 switches (this is documented as a necessary step in the procedures in [Appendix C](#). See the *SANbox-8/16 Switch Management User's Manual* for instructions on removing the preconfigured hard zones on all Sun StorEdge Network FC Switch-8 and Switch-16 switches.

Known Documentation Problems

This section discusses documentation errors you might encounter and steps to correct these problems. This information is in addition to known documentation problems documented in the *Sun Cluster 3.0 12/01 Release Notes*.

Hardware Guide

Sun Netra st D1000

The Netra™ st D1000 has been qualified for use with Sun Cluster 3.0, but maintenance procedures do not yet exist in the *Sun Cluster 3.0 12/01 Hardware Guide* specifically for the Netra st D1000. However, you can use the existing *Sun Cluster 3.0 12/01 Hardware Guide* procedures for the StorEdge D1000 to maintain a Netra st D1000 in a cluster environment. Substitute the following document references in place of the StorEdge D1000 document references if you are maintaining a Netra st D1000.

- *Netra st A1000 and Netra st D1000 Installation and Maintenance Manual* (805-7147)
- *Netra st A1000 and Netra st D1000 Product Notes* (805-7148)

Sun StorEdge T3/T3+ (Single-Controller Configuration)

The *Sun Cluster 3.0 12/01 Hardware Guide* chapter on “Installing and Maintaining a Sun StorEdge T3 or T3+ Array Single-Controller Configuration” incorrectly indicates that the procedure for upgrading a Sun StorEdge T3 array controller to a Sun StorEdge T3+ array controller does not require any cluster-specific steps. The correct procedure should be as follows.

How To Upgrade a StorEdge T3 Controller to a StorEdge T3+ Controller



Caution – Perform this procedure on one array at a time. This procedure requires that you take the array in which you are upgrading the controller offline. If you take more than one array offline, your cluster will lose access to data, if the arrays are submirrors of each other.

- 1. On one node attached to the StorEdge T3 array in which you are upgrading the controller, detach that array's submirrors.**

For more information, see your Solstice DiskSuite™ or VERITAS Volume Manager documentation.

- 2. Upgrade the StorEdge T3 array controller to a StorEdge T3+ array controller.**

See the *Sun StorEdge T3 Array Controller Upgrade Manual* for instructions.

- 3. Reattach the submirrors to resynchronize them.**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

Generic Data Services

This appendix provides information on the generic data service (GDS) and shows you how to create a service that uses GDS using either the command line interface or the SunPlex™ Agent Builder.

- [“Generic Data Service Overview” on page 11](#)
- [“Using the Command Line to Create a Service Using GDS” on page 16](#)
- [“Using the SunPlex Agent Builder to Create a Service Using GDS” on page 19](#)

Generic Data Service Overview

A generic data service (GDS) is a mechanism for making simple applications highly available or scalable by plugging them into the Sun Cluster's Resource Group Management framework. This mechanism does not require the coding of an agent which is the typical approach for making an application highly available or scalable.

The Generic Data Service (GDS) is a single, precompiled data service that you can use to make various applications highly available or scalable. In this approach, the precompiled data service and its components, the callback method implementations and the RTR file, are not subject to change.

The GDS model is especially attractive to people who develop custom HA solutions. With the GDS model, developers are able to separate the user modifiable components from the unchangeable ones and thus are able to better diagnose problems on a customer site.

Precompiled Resource Type

The generic data service resource type, `SUNW.gds`, is included in the `SUNWscgds` package. The `scinstall` utility installs this package on the cluster during cluster installation. The `SUNWscgds` package includes the following files:

```
# pkgchk -v SUNWscgds

/opt/SUNWscgds
/opt/SUNWscgds/bin
/opt/SUNWscgds/bin/gds_monitor_check
/opt/SUNWscgds/bin/gds_monitor_start
/opt/SUNWscgds/bin/gds_monitor_stop
/opt/SUNWscgds/bin/gds_probe
/opt/SUNWscgds/bin/gds_svc_start
/opt/SUNWscgds/bin/gds_svc_stop
/opt/SUNWscgds/bin/gds_update
/opt/SUNWscgds/bin/gds_validate
/opt/SUNWscgds/etc
/opt/SUNWscgds/etc/SUNW.gds
```

Methods of Creating a Service Using GDS

There are 2 methods of creating a service that uses GDS:

- GDS using the command line interface
- GDS using the SunPlex Agent Builder

GDS and the Command Line Interface

This method utilizes the precompiled data service code in `SUNWscgds` but requires that the developer use the command line interface (`scrgadm` and `scswitch`) to enter the parameters required to create and configure the resource.

GDS and the SunPlex Agent Builder

Use the SunPlex Agent Builder and select GDS to generate the driving scripts which rely on the precompiled code being available in `SUNWscgds`. The user input is used to generate a set of driving scripts that configure resources for the given application.

How to Select Which Method to Use to Create a GDS

As shown in the procedures [“Command Line Interface to Create a Highly Available Service Using GDS” on page 16](#) and [“Command Line Interface to Create a Scalable Service Using GDS” on page 17](#), a significant amount of typing is required to issue the appropriate `scradm` commands.

Using GDS with SunPlex Agent Builder simplifies the process since the precompiled binaries are used and there is significant less typing of input information.

Mandatory Inputs

The start command launches the application. It should be a UNIX command complete with its arguments that is acceptable by `pmfadm`. For network aware (or client-server) applications, the command to launch the application may contain the `hostname` (or IP address) it's going to listen on. In this case, the user has to configure the `LogicalHostname` or `SharedAddress` resources separately. The GDS start method launches the application under the control of Process Monitoring Facility (PMF) using the start command.

Optional Inputs

The optional inputs to GDS include the following:

- Stop command
- Probe command
- Timeout for start method
- Timeout for stop method
- Timeout for fault monitor probe
- Child monitor level (only with GDS using command line interface)

Stop Command

The stop command is supposed to stop the application and not return until the application has been completely stopped. It should be a UNIX command complete with its arguments that is acceptable by `hatimerun`.

If the stop command is provided, the GDS stop method launches the stop command with 80% of the timeout. Regardless of the outcome of launching the stop command, the GDS stop method still sends `SIGKILL` with 15% of the stop timeout. The remaining 5% of the time is assumed to have been used by the house keeping overhead.

If the stop command is omitted, the GDS tries to stop the application using signals. GDS first sends `SIGTERM` with 80% of the stop timeout. If the `SIGTERM` fails to stop the application, GDS then sends `SIGKILL` with 15% of the stop timeout. The remaining 5% of the time is assumed to have been used by the house keeping overhead.

Probe Command

The probe command checks the health of the given application. It should be a UNIX command complete with its arguments that is acceptable by `hatimerun`. The exit status of the probe command is used to determine the severity of the failure of the application. This exit status, called probe status, must be an integer between 0 (for success) and 100 (for complete failure). The probe status is used with in the GDS probing algorithm to take the decision of restarting the application locally, versus failing it over to another node.

If the probe command is omitted and the application uses network, GDS provides its own simple probe. This alternative probe simply connects to the application on the given `hostname` and port number. If the connect succeeds, it disconnects immediately. If both connect and disconnect succeed, the application is deemed to be running healthily.

Note – The probe provided with GDS is only intended to be a simple substitute for the fully functioning application specific probe.

Start Method Timeout

The default timeout for the start method is 300 seconds.

Stop Method Timeout

The default timeout for the stop method is 300 seconds.

Fault Monitor Probe Timeout

The default timeout for the fault monitor probe is 30 seconds.

Child Monitor Level

This option provides more control over what level of processes get monitored through PMF. It denotes the level up to which the forked children processes have to be monitored. This is the `-C` argument to the `pmfadm` command. See the `pmfadm(1M)` man page for more details.

Note – This option can only be specified using the command line interface. This option cannot be specified if you are using the SunPlex Agent Builder.

Using the Command Line to Create a Service Using GDS

In this section we describe how the above mentioned parameters would actually be input to the GDS. The GDS will be used and administered using the existing Sun Cluster administrative commands, such as `scrgadm(1M)` and `scswitch(1M)`.

▼ Command Line Interface to Create a Highly Available Service Using GDS

1. Register the resource type `SUNW.gds`.

```
# scrgadm -a -t gds
```

2. Create the resource group containing the LogicalHostname resource and the failover service itself.

```
# scrgadm -a -g haapp_rg
```

3. Create the resource for the LogicalHostname resource.

```
# scrgadm -a -L -g haapp_rs -l hhead
```

4. Create the resource for the failover service itself.

```
# scrgadm -a -j haapp_rs -g haapp_rg -t SUNW.gds \  
-y Scalable=false -y Start_timeout=120 \  
-y Stop_timeout=120 -x Probe_timeout=120 \  
-x Start_command="/usr/local/ha/appctl start" \  
-x Stop_command="/usr/local/ha/appctl stop" \  
-x Probe_command="/usr/local/app/bin/probe" \  
-x Child_mon_level=0 -y Network_resources_used=hhead
```

Note – See [“Start Command Arguments” on page 18](#) for information on issuing the start command.

5. Bring the resource group `haapp_rg` online.

```
# scswitch -Z -g haapp_rg
```

▼ Command Line Interface to Create a Scalable Service Using GDS

1. Register the resource type `SUNW.gds`.

```
# scrgadm -a -t gds
```

2. Create the resource group for the SharedAddress resource.

```
# scrgadm -a -g sa_rg
```

3. Create the SharedAddress resource on `sa_rg`.

```
# scrgadm -a -S -g sa_rg -l hhead
```

4. Create the resource group for the scalable service.

```
# scrgadm -a -g app_rg -y Maximum primaries=2 \  
-y Desired primaries=2 -y RG_dependencies=sa_rg
```

5. Create the resource group for the scalable service itself.

```
# scrgadm -a -j app_rs -g app_rg -t SUNW.gds \  
-y Scalable=true -y Start_timeout=120 \  
-y Stop_timeout=120 -x Probe_timeout=120 \  
-x Start_command="/usr/local/app/bin/start" \  
-x Stop_command="/usr/local/app/bin/stop" \  
-x Probe_command="/usr/local/app/bin/probe" \  
-x Child_mon_level=0 -y Network_resource_used=hhead
```

Note – See [“Start Command Arguments”](#) on page 18 for information on issuing the start command.

6. Bring the resource group containing the network resources online.

```
# scswitch -Z -g sa_rg
```

7. Bringing the resource group `app_rg` online.

```
# scswitch -Z -g app_rg
```

Start Command Arguments

The arguments to the start script will vary according to the type of resource. Check the customized man page or run the start script without any argument to get a usage statement.

```
# /opt/NETapp/util/startapp
```

The resource name of LogicalHostname or SharedAddress must be specified.

For failover services:

Usage: `startapp -h <logical host name>`

`-p <port and protocol list>`

`[-n <nafogroup/adapter list>]`

`[-v]`

For scalable services:

Usage: `startapp`

`-h <shared address name>`

`-p <port and protocol list>`

`[-l <load balancing policy>]`

`[-n <nafogroup/adapter list>]`

`[-w <load balancing weights>] [-v]`

For stand-alone (non-network aware applications):

Usage: `startapp [-v]`

Using the SunPlex Agent Builder to Create a Service Using GDS

You can use the SunPlex Agent Builder to create the service that uses GDS. See the *Sun Cluster 3.0 12/01 Data Services Developer's Guide* for information on the SunPlex Agent Builder as well as a definition of the fields.

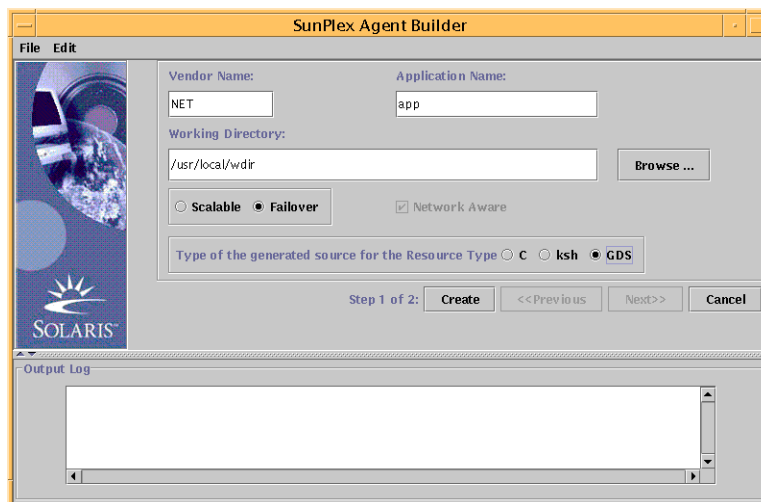
Create a Service Using GDS in the SunPlex Agent Builder

▼ Creating a Service Using GDS in the Agent Builder

1. Bring up the SunPlex Agent Builder.

```
# /usr/cluster/bin/scdsbuilder
```

2. The SunPlex Agent Builder panel appears.

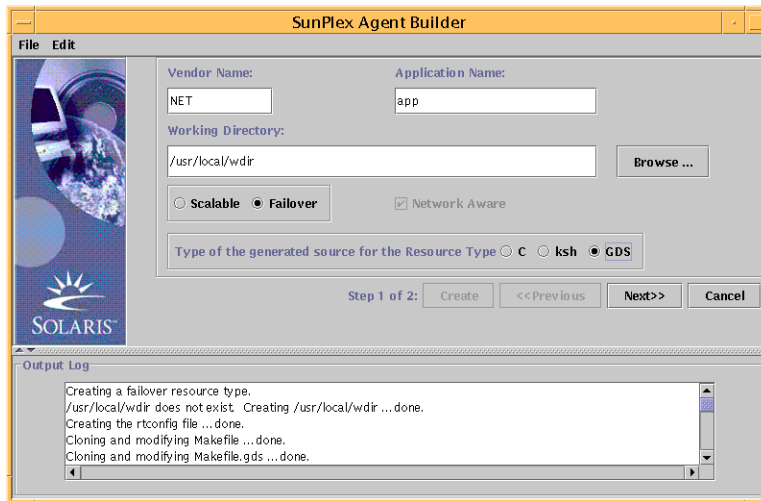


3. Enter the Vendor Name.
4. Enter the Application Name.

Note – The combination of Vendor and Application Name may not contain more than nine characters. It will be used as the name of the package for the driving scripts.

5. Enter the Working Directory. You can use the Browse pull-down to select the directory rather than typing the path.
6. Select whether the data service is scalable or failover.
7. You do not need to select Network Aware since that is the default when creating GDS.
8. Select GDS.
9. Click the Create button to create the driving scripts.

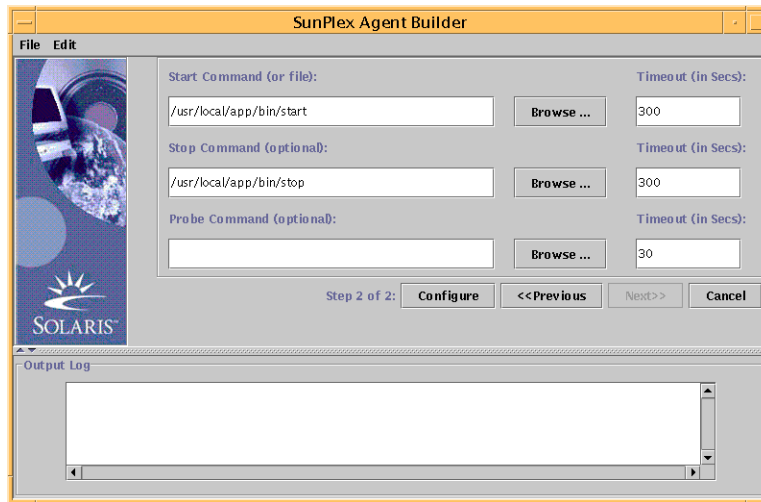
The SunPlex Agent Builder panel displays the results of the creation of the service. The Create button is grayed out and the Next button is now available.



▼ Configuring the Driving Scripts of the Service

After creating the driving scripts, you need to use the SunPlex Agent Builder to configure the new service.

1. Click the Next button and the configuration panel appears.



2. Either enter the location of the Start command or use the browse button to locate the Start command.
3. (Optional) Enter the Stop command or use the browse button to locate the Stop command.
4. (Optional) Enter the Probe command or use the browse button to locate the Probe command.
5. (Optional) Specify the timeout values for the Start, Stop, and Probe commands.
6. Click Configure to start the configuration of the driving scripts.

Note – The package name will consist of a concatenation of the Vendor Name and Application Name.

A package for driving scripts will be created and placed in the following file:

`<working-dir>/<vendor_name><application>/pkg`

For example, `/usr/local/wdir/NETapp/pkg`.

7. Install the completed package on both nodes of the cluster.

```
# cd /usr/local/wdir/NETapp/pkg
# pkgadd -d . NETapp
```

8. The following files will be installed during the pkgadd:

```
/opt/NETapp
/opt/NETapp/README.app
/opt/NETapp/man
/opt/NETapp/man/man1m
/opt/NETapp/man/man1m/removeapp.1m
/opt/NETapp/man/man1m/startapp.1m
/opt/NETapp/man/man1m/stopapp.1m
/opt/NETapp/man/man1m/app_config.1m
/opt/NETapp/util
/opt/NETapp/util/removeapp
/opt/NETapp/util/startapp
/opt/NETapp/util/stopapp
/opt/NETapp/util/app_config
```

Note – The man pages and script names will correspond to the Application Name you entered above preceded by the script name; for example, startappname.

To view the man pages, you need to specify the path to the man page. For example, to view the startapp man pages, use the following command:

```
% man -M /opt/NETapp/man startapp
```

9. On one node of the cluster, configure the resources and start the application.

```
# /opt/NETapp/util/startapp -h logichostname \  
-p port-and-protocol-list
```

The arguments to the start script will vary according to the type of resource: failover or scalable. Check the customized man page or run the start script without any argument to get a usage statement.

```
# /opt/NETapp/util/startapp  
The resource name of LogicalHostname or SharedAddress must be  
specified.
```

For failover services:
Usage: startapp -h *<logical host name>*
 -p *<port and protocol list>*
 [-n *<nafogroup/adaptor list>*]
 [-v]

For scalable services:
Usage: startapp
 -h *<shared address name>*
 -p *<port and protocol list>*
 [-l *<load balancing policy>*]
 [-n *<nafogroup/adaptor list>*]
 [-w *<load balancing weights>*] [-v]

For stand-alone (non-network aware applications):
Usage: startapp [-v]

Output from SunPlex Agent Builder

The SunPlex Agent Builder generates three driving scripts based on input you provide during package creation as well as a configuration file. The configuration files specify the names of the resource group and resource type.

Customizable Driving Scripts

The following are the driving scripts:

- Start script: Used to configure the resources and start the application under cluster control
- Stop script: Used to stop the application and take down resources and resource groups

- Remove script: Used to remove the resources and resource groups created by the start script

You can customize these scripts to provide your own names for resource groups or other parameters that are normally given as inputs to the `scrgadm` command. If you do not customize the scripts, the driving scripts will provide reasonable defaults for the `scrgadm` parameters.

Configuration File for an HA Service

Following is a sample configuration file for a highly available service. The default name for the resource group and resource application are `applicationname-hars` and `applicationname-harg` where `applicationname` is the name of the application that you specified to the SunPlex Agent Builder.

You can modify these names if you want to use an existing resource group. You must make any required changes before you run the start script. Open the configuration file using your favorite editor and make the changes.

```
#
# Parameter file for application generated by SunPlex Data Service
# Builder
#
# This file will be sourced in by startapp, stopapp and removeapp
# The parameters listed below will be used in those scripts.
# These parameters can be customized in (key=value) form
#
# RS_NAME:      name of the resource for the application
# RG_NAME:      name of the resource group containing RS_NAME
# MAXIMUM_PRIMARYES: Maximum primaries of RG_NAME
# DESIRED_PRIMARYES: Desired primaries of RG_NAME
#
# For example:
#
# RS_NAME=rs_app
# RG_NAME=rg_app
# MAXIMUM_PRIMARYES=1
# DESIRED_PRIMARYES=1
#
RS_NAME=app-hars
RG_NAME=app-harg
MAXIMUM_PRIMARYES=1
DESIRED_PRIMARYES=1
```

Configuration File for Scalable Services

Following is a sample configuration file for a scalable service. The default names for the resource application and resource group are *applicationname-sars* and *applicationname-sarg* where *applicationname* is the name of the application that you specified to the SunPlex Agent Builder. The name of the SharedAddress resource is the concatenation of *sa-rg-* and the name of the *SharedAddress*.

You can modify these names if you want to use an existing resource group or SharedAddress. You must make any required changes before you run the start script.

Open the configuration file using your favorite editor and make any required changes.

```
#
# Parameter file for application generated by SunPlex Data Service
# Builder
#
# This file will be sourced in by startecho, stopecho and removeecho
# The parameters listed below will be used in those scripts.
# These parameters can be customized in (key=value) form
#
# RS_NAME:      name of the resource for the application
# RG_NAME:      name of the resource group containing RS_NAME
# SA_RG_NAME:   the resource group containing the Shared Address
resource
# MAXIMUM_PRIMARIES: Maximum primaries of RG_NAME
# DESIRED_PRIMARIES: Desired primaries of RG_NAME
#
# For example:
#
# RS_NAME=rs_echo
# RG_NAME=rg_echo
# SA_RG_NAME=sa_rg
# MAXIMUM_PRIMARIES=4
# DESIRED_PRIMARIES=4
#
RS_NAME=echo-sars
RG_NAME=echo-sarg
SA_RG_NAME=sa-rg-`${HOSTNAME}`
MAXIMUM_PRIMARIES=${NO_OF_NODES}
DESIRED_PRIMARIES=${NO_OF_NODES}
```

Command Line Equivalent of the SunPlex Agent Builder

You can use the command line interface to create a service that uses DGS that is the equivalent of using the SunPlex Agent Builder.

▼ Creating A Service that Uses GDS Using the Command Line Equivalent of the SunPlex Agent Builder

1. Create the service.

For a failover service, use the following command:

```
# sdcscrcate -g -V NET -T app -d /usr/local/wdir
```

For a scalable service, use the following command:

```
# sdcscrcate -g -s -V NET -T app -d /usr/local/wdir
```

Note – The *-d* parameters are optional. If it is not specified, the working directory defaults to the current directory.

2. Configure the service.

```
# sdcscconfig -s "/usr/local/app/bin/start" \  
-t "/usr/local/app/bin/stop" \  
-m "/usr/local/app/bin/probe" \  
-d /usr/local/wdir
```

Note – Only the start command is required. All other parameters are optional.

3. Install the completed package on both nodes of the cluster.

```
# cd /usr/local/wdir/NETapp/pkg  
# pkgadd -d . NETapp
```

4. The following files will be installed during the pkgadd:

```
/opt/NETapp  
/opt/NETapp/README.app  
/opt/NETapp/man  
/opt/NETapp/man/man1m  
/opt/NETapp/man/man1m/removeapp.1m  
/opt/NETapp/man/man1m/startapp.1m  
/opt/NETapp/man/man1m/stopapp.1m  
/opt/NETapp/man/man1m/app_config.1m  
/opt/NETapp/util  
/opt/NETapp/util/removeapp  
/opt/NETapp/util/startapp  
/opt/NETapp/util/stopapp  
/opt/NETapp/util/app_config
```

Note – The man pages and script names will correspond to the Application Name you entered above preceded by the script name; for example, startappname.

To view the man pages, you need to specify the path to the man page. For example, to view the `startapp` man pages, use the following command:

```
% man -M /opt/NETapp/man startapp
```

5. On one node of the cluster, configure the resources and start the application.

```
# /opt/NETapp/util/startapp -h logicalhostname \  
-p port-and-protocol-list
```

The arguments to the start script will vary according to the type of resource: failover or scalable. Check the customized man page or run the start script without any argument to get a usage statement.

```
# /opt/NETapp/util/startapp  
The resource name of LogicalHostname or SharedAddress must be  
specified.
```

```
For failover services:  
Usage: startapp -h logical host name  
       -p port and protocol list  
       [-n nafogroup/adaptor list]  
       [-v]
```

```
For scalable services:  
Usage: startapp  
       -h shared address name  
       -p port and protocol list  
       [-l load balancing policy]  
       [-n nafogroup/adaptor list]  
       [-w load balancing weights] [-v]
```

```
For stand-alone (non-network aware applications):  
Usage: startapp [-v]
```

Installing and Configuring Sun Cluster HA for SAP

This chapter provides instructions on how to plan, set up, and configure Sun Cluster HA for SAP on your Sun Cluster nodes.

This chapter includes the following procedures.

- [“How to Upgrade a Resource Type or Convert a Failover Application Resource to a Scalable Application Resource” on page 39](#)
- [“How to Install SAP and Database” on page 41](#)
- [“How to Install SAP Scalable Application Server” on page 42](#)
- [“How to Enable SAP to Run in the Cluster” on page 44](#)
- [“How to Verify SAP and the Database Installation with Central Instance” on page 45](#)
- [“How to Verify SAP and the Database Installation with Failover Application Server” on page 48](#)
- [“How to Register and Configure Sun Cluster HA for SAP Central Instance” on page 51](#)
- [“How to Register and Configure Sun Cluster HA for SAP with Failover Application Server” on page 52](#)
- [“How to Register and Configure Sun Cluster HA for SAP with Scalable Application Server” on page 53](#)
- [“How to Verify the Installation and Configuration of Sun Cluster HA for SAP, Central Instance, and DBMS” on page 55](#)
- [“How to Verify the Installation and Configuration of Sun Cluster HA for SAP with Failover Application Server” on page 56](#)
- [“How to Verify the Installation and Configuration of Sun Cluster HA for SAP with Scalable Application Server” on page 57](#)
- [“How to Set Up a Lock File for Central Instance” on page 63](#)
- [“How to Set Up a Lock File for Scalable Application Server” on page 64](#)

Sun Cluster HA for SAP Overview

Sun Cluster HA for SAP provides fault monitoring and automatic failover for the SAP application to eliminate single points of failure in an SAP system. The following table lists the data services that best protect SAP components in a Sun Cluster configuration. You can configure Sun Cluster HA for SAP as a failover application or a scalable application.

For conceptual information on failover and scalable services, see the *Sun Cluster 3.0 12/01 Concepts*.

TABLE B-1 Protection of SAP Components

SAP Component	Protected by
SAP database	Sun Cluster HA for Oracle, if the database is Oracle
SAP central instance	Sun Cluster HA for SAP, the resource type is <code>SUNW.sap_ci</code> or <code>SUNW.sap_ci_v2</code>
SAP application server	Sun Cluster HA for SAP, the resource type is <code>SUNW.sap_as</code> or <code>SUNW.sap_as_v2</code>
NFS file system	Sun Cluster HA for NFS

Use the `scinstall(1M)` command to install Sun Cluster HA for SAP. Sun Cluster HA for SAP requires a functioning cluster with the initial cluster framework already installed. See the *Sun Cluster 3.0 12/01 Software Installation Guide* for details on initial installation of clusters and data service software. Register Sun Cluster HA for SAP after you successfully install the basic components of the Sun Cluster and SAP software.

Installing and Configuring Sun Cluster HA for SAP

The following table lists the sections that describe the installation and configuration tasks.

TABLE B-2 Task Map: Installing and Configuring Sun Cluster HA for SAP

Task	For Instructions, Go To
Plan the SAP installation	“Sun Cluster HA for SAP Overview” on page 30
	“Configuration Guidelines for Sun Cluster HA for SAP” on page 34
	“Sample Configurations” on page 35
	“Pre-Installation Requirements” on page 37
	“Pre-Installation Considerations” on page 38
Upgrade Sun Cluster HA for SAP	“How to Upgrade a Resource Type or Convert a Failover Application Resource to a Scalable Application Resource” on page 39
Install and configure SAP and the database	“How to Install SAP and Database” on page 41
	“How to Install SAP Scalable Application Server” on page 42
	“How to Enable SAP to Run in the Cluster” on page 44
	“How to Verify SAP and the Database Installation with Central Instance” on page 45
	“How to Verify SAP and the Database Installation with Failover Application Server” on page 48
Configure the Sun Cluster HA for DBMS	“Configuring Sun Cluster HA for DBMS” on page 50

TABLE B-2 Task Map: Installing and Configuring Sun Cluster HA for SAP

Task	For Instructions, Go To
Configure Sun Cluster HA for SAP with a failover application server	“How to Register and Configure Sun Cluster HA for SAP Central Instance” on page 51 “How to Register and Configure Sun Cluster HA for SAP with Failover Application Server” on page 52
or	
Configure with a scalable application server	“How to Register and Configure Sun Cluster HA for SAP Central Instance” on page 51 “How to Register and Configure Sun Cluster HA for SAP with Scalable Application Server” on page 53
Verify Sun Cluster HA for SAP installation and configuration	“How to Verify the Installation and Configuration of Sun Cluster HA for SAP, Central Instance, and DBMS” on page 55 “How to Verify the Installation and Configuration of Sun Cluster HA for SAP with Failover Application Server” on page 56 “How to Verify the Installation and Configuration of Sun Cluster HA for SAP with Scalable Application Server” on page 57
Configure SAP extension properties	“Configuring Sun Cluster HA for SAP Extension Properties” on page 58
Understand Sun Cluster HA for SAP fault monitor	“Sun Cluster HA for SAP Fault Monitor” on page 66

Planning the Installation and Configuration

Read the following information to plan your Sun Cluster HA for SAP installation.

- [Sun Cluster HA for SAP Packages and Support](#)
- [Configuration Guidelines for Sun Cluster HA for SAP](#)
- [Sample Configurations](#)
- [Pre-Installation Considerations](#)

Sun Cluster HA for SAP Packages and Support

The following tables lists the packages that Sun Cluster HA for SAP supports.

TABLE B-3 Sun Cluster HA for SAP packages from Sun Cluster 3.0 7/01

Resource Type	Description
SUNW.sap_ci	Added support for failover central instance.
SUNW.sap_as	Added support for failover application servers.

TABLE B-4 Sun Cluster HA for SAP package from Sun Cluster 3.0 12/01

Resource Type	Description
SUNW.sap_ci	Same as Sun Cluster 3.0 7/01. See TABLE B-3 .
SUNW.sap_as	Same as Sun Cluster 3.0 7/01. See TABLE B-3 .
SUNW.sap_ci_v2	Added the Property <code>Network_resources_used</code> resource property to the Resource Type Registration (RTR) file. Retained support for failover central instance.
SUNW.sap_as_v2	Added the Property <code>Network_resources_used</code> resource property to RTR file. Added support for scalable application servers. Retained support for failover application servers

Configuration Guidelines for Sun Cluster HA for SAP

When you design a Sun Cluster HA for SAP configuration, consider the following guidelines.

Failover and Scalable Applications

- **Use an SAP software version that is qualified with Sun Cluster 3.0.**
- **Use an SAP software version with automatic enqueue reconnect mechanism capability** – Sun Cluster HA for SAP relies on this capability. SAP 4.0 software with patch information and later releases should have automatic enqueue reconnect mechanism capability.
- **Retrieve the latest patch for the `sapstart` executable** – This patch enables Sun Cluster HA for SAP users to configure a lock file. For details on the benefits of this patch in your cluster environment, see [“Setting Up a Lock File” on page 63](#).
- **Read all of the related SAP online service-system notes for the SAP software release and database that you are installing on your Sun Cluster configuration** – Identify any known installation problems and fixes.
- **Consult SAP software documentation for memory and swap recommendations** – SAP software uses a large amount of memory and swap space.
- **Generously estimate the total possible load on nodes that might host the central instance, the database instance, and the application server, if you have an internal application server** – This guideline is especially important if you configure the cluster to ensure that the central instance, database instance, and application server will all exist on one node if failover occurs.
- **Install application servers on either the same cluster that hosts the central instance or on a separate cluster** – If you install and configure any application server outside of the cluster environment, Sun Cluster HA for SAP does not perform fault monitoring and does not automatically restart or fail over those application servers. You must manually start and shut down application servers that you install and configure outside of the cluster environment.
- **Limit node names and network resource names to ten characters or less** – This limitation is an SAP software requirement.

Scalable Applications

- **Use the same instance number and the same SID when you create all application server instances on multiple cluster nodes** - This guideline ensures ease of maintenance and ease of administration because you will only need to use one set of commands to maintain all application servers on multiple nodes.
- **Install the application servers locally on the cluster node instead of on a cluster file system** - This guideline ensures that another application server does not overwrite the `log/data/work/sec` directory for the application server.
- **Ensure that the `SAPSIDadm` home directory resides on a cluster file system** - This guideline enables you to maintain only one set of scripts for all application server instances that run on all nodes. However, if you have some application servers that need to be configured differently (for example, application servers with different profiles), install those application servers with different instance numbers, and then configure them in a separate resource group.
- **Place the application servers into multiple resource groups if you want to use RGOffload functionality to shut down one or more application servers when a higher priority resource is failing over** - This guideline provides flexibility and availability if you want to use RGOffload functionality (a separate resource type) to offload one or more application servers for the database. The functionality you gain from this guideline supersedes the ease of use you gain from placing the application servers into one large group. See “Freeing Node Resources by Offloading Non-critical Resource Groups” in the *Sun Cluster 3.0 12/01 Data Services Installation and Configuration Guide* for more information on using the RGOffload resource type.

Sample Configurations

See your Enterprise Services representative for the most current information about supported SAP versions. The following figures illustrate sample configurations for Sun Cluster HA for SAP.




FIGURE B-1 Four-Node Cluster with Central Instance, Application Servers, and Database




FIGURE B-2 Two-Node Cluster with Central Instance, NFS, and Non-HA External Application Servers

Note – The configuration in [FIGURE B-2](#) was a common configuration under previous Sun Cluster releases. To use the Sun Cluster 3.0 software to the fullest extent, configure SAP as shown in [FIGURE B-1](#) or [FIGURE B-3](#).




FIGURE B-3 Two-Node Cluster With Central Instance and Development Node

Pre-Installation Requirements

Before you install the SAP software, see [“Installing and Configuring SAP and Database” on page 41](#), understand the following requirements.

- **After you create all of the file systems for the database and for SAP software, create the mount points, and put the mount points in the `/etc/vfstab` file on all of the cluster nodes** – See the SAP installation guides, *Installation of the SAP R/3 on UNIX* and *R/3 Installation on UNIX-OS Dependencies*, for details on how to set up the database and SAP file systems.
- **Create the required groups and users on all of the cluster nodes** – See the SAP installation guides, *Installation of the SAP R/3 on UNIX* and *R/3 Installation on UNIX-OS Dependencies*, for details on how to create SAP groups and users.
- **Configure Sun Cluster HA for NFS on the cluster that hosts the central instance if you plan to install some external SAP application servers** – See “Installing and Configuring Sun Cluster HA for NFS” in the *Sun Cluster 3.0 12/01 Data Services Installation and Configuration Guide* for details on how to configure Sun Cluster HA for NFS.
- **Set up the `/etc/nsswitch.conf` file so that the data service starts and stops correctly during switchovers or failovers** – On each node that can master the logical host that runs Sun Cluster HA for SAP, the `/etc/nsswitch.conf` file must have one of the following entries for `group`.

```
group:  
group: files  
group: files [NOTFOUND=return] nis  
group: files [NOTFOUND=return] nisplus
```

Sun Cluster HA for SAP uses the `su user` command to start and stop the database node. The network information name service might become unavailable when a cluster node’s public network fails. Add one of the preceding entries for `group` to ensure that the `su(1M)` command does not refer to the NIS/NIS+ name services if the name service becomes unavailable.

Pre-Installation Considerations

Before you install the SAP software, see [“Installing and Configuring SAP and Database” on page 41](#), install SAP binaries and SAP users’ home directories. Install SAP binaries and users’ home directories on a cluster file system. Installation on a cluster file system, however, has some drawbacks with SAP software release upgrades. See “Determining the Location of the Application Binaries” in the *Sun Cluster 3.0 12/01 Data Services Installation and Configuration Guide* for information about drawbacks.

Upgrading Sun Cluster HA for SAP

As [TABLE B-3](#) and [TABLE B-4](#) illustrate, the Sun Cluster HA for SAP package from Sun Cluster 3.0 7/01 does not support a scalable application server and the `Network_resources_used` resource property. Therefore, you have the following upgrade options.

- Retain (do not upgrade) the existing `SUNW.sap_ci` and `SUNW.sap_as` resource types. Choose this option if any of the following statements apply to you.
 - You cannot schedule down time.
 - You do not want the `Network_resources_used` resource property.
 - You do not want to configure a scalable application server.
- Upgrade a resource type.

See [“How to Upgrade a Resource Type or Convert a Failover Application Resource to a Scalable Application Resource” on page 39](#) for the procedure on how to upgrade a resource type.
- Convert a failover application resource to a scalable application resource.

See [“How to Upgrade a Resource Type or Convert a Failover Application Resource to a Scalable Application Resource” on page 39](#) for the procedure on how to convert a failover application resource to a scalable application resource.

▼ How to Upgrade a Resource Type or Convert a Failover Application Resource to a Scalable Application Resource

To upgrade a resource type or to convert a failover application server resource to a scalable application server resource, perform the following steps. This procedure requires that you schedule down time.

1. **Disable the existing resource.**
2. **Delete the existing resource from the resource group.**
3. **Delete the existing resource type if no other resource uses it.**
4. **Which task are you performing?**
 - If you are upgrading a resource type, skip to [Step 6](#).
 - If you are converting a failover application server resource to a scalable application server resource, proceed to [Step 5](#).

5. **Create the new application server resource group `scalable_rg` if you do not want to create the scalable application server resource in an existing scalable server resource group.**
6. **Create the new resource in the resource group.**

Where to Go From Here

Go to [“How to Register and Configure Sun Cluster HA for SAP with Scalable Application Server”](#) on page 53.

Installing and Configuring SAP and Database

Use the procedures in this section to perform the following tasks.

- Install SAP and the database.
- Install SAP and the scalable application server.
- Enable SAP to run in the cluster.
- Verify SAP and the database installation with the central instance.
- Verify SAP and the database installation with the failover application server.

▼ How to Install SAP and Database

To install SAP and the database, perform the following steps.

1. **Become superuser on one of the nodes in the cluster where you are installing the central instance.**
2. **Install SAP binaries on a cluster file system.**

Note – Before you install SAP software on a cluster file system, use the `scstat(1M)` command to verify that the Sun Cluster software is fully operational.

- a. **For all of the SAP-required kernel parameter changes, edit the `/etc/system` file on all of the cluster nodes that will run the SAP application.**

After you edit the `/etc/system` file, reboot each node. See the SAP document *R/3 Installation on UNIX-OS Dependencies* for details on kernel parameter changes.

- b. **See the SAP document *Installation of the SAP R/3 on UNIX* for details on how to install the central instance, the database, and the application server instances.**

See [“How to Install SAP Scalable Application Server” on page 42](#) for the procedure on how to install a scalable application server in a Sun Cluster environment.

Where to Go From Here

Go to [“How to Enable SAP to Run in the Cluster” on page 44](#).

▼ How to Install SAP Scalable Application Server

To install scalable application server instances, perform the following steps. This procedure assumes that you installed the central instance and the database.

Tip – The following file system layout ensures ease of use and prevents data from being overwritten.

- Cluster File Systems

 - `/sapmnt/SID`

 - `/usr/sap/SID` -> all subdirectories except the *app-instance* subdirectory

 - `/usr/sap/SID/home` -> the SAPSIDadm home directory

 - `/usr/sap/trans`

- Local File Systems

 - `/usr/sap/local/SID/app-instance`

1. Create all SAP directories on cluster file systems.

- Ensure that the central instance and the database can fail over.
- Set up the lock file on cluster file system for the central instance to prevent a multiple startup from a different node.

For the procedure on how to set up a lock file on the central instance, see [“How to Set Up a Lock File for Central Instance” on page 63](#).

- Ensure that all application servers can use the SAP binaries on a cluster file system.

2. Install the central instance and the database on a cluster file system.

See the SAP document *Installation of the SAP R/3 on UNIX* for details on how to install the central instance and the database.

3. On all nodes that will host the scalable application server, create a local directory for the `data/log/sec/work` directories and the log files for starting and stopping the application server.

Create a local directory for each new application server.



Caution – You must perform this step. If you do not perform this step, you will inadvertently install a different application server instance on a cluster file system and the two application servers will overwrite each other.

Example:

```
# mkdir -p /usr/sap/local/SC3/D03
```

4. Set up a link to point to the local application server directory from a cluster file system, so the application server and the `startup/stop` log file will be installed on the local file system.

```
# ln -s /usr/sap/local/SC3/D03 /usr/sap/SC3/D03
```

5. Install the application server.
6. Edit the `start/stop` script so that the `startup/stop` log files will be node specific under `users sapsidadm` and `orasapsid`.

Example:

```
# vi startsap_D03
```

Before

```
LOGFILE=$R3S_LOGDIR/`basename $0`.log
```

After

```
LOGFILE=$R3S_LOGDIR/`basename $0`_'uname -n'.log
```

7. Copy the application server (with the same SAPSID and the same instance number) on all nodes that run the scalable application server.

The nodes that run the scalable application server are in the scalable application server resource group `nodelist`.

8. Ensure that you can startup and stop the application server from each node, and verify that the log files are in the correct location.
9. Install the SAP logon group if you use a logon group.

Where to Go From Here

Go to [“How to Enable SAP to Run in the Cluster”](#) on page 44.

▼ How to Enable SAP to Run in the Cluster

During SAP installation, the SAP software creates files and shell scripts on the server on which you installed the SAP central instance. These files and scripts use physical server names. To run the SAP software with Sun Cluster software, replace references to a physical server with references to a network resource (logical hostname). Throughout these steps, the term *physicalserver* represents a physical server, and the term *logical-hostname* represents a network resource.

To enable SAP to run in the cluster, perform the following steps.

1. **Make backup copies of the files that you will modify in the following steps.**
2. **Log in to the node on which you installed the SAP software.**
3. **Shut down all SAP instances (central instance and application server instances) and the database.**

Note – In addition to the central instance and database, shut down any application servers that are running.

4. **Become user *sapsidadm*, and then perform the following tasks.**
 - In the SPSIDadm home directory, modify all of the file names that reference a physical server name.
 - In the SPSIDadm home directory, modify all of the file contents—except log file contents—that reference a physical server name.
 - In the SAP profile directory, modify all of the file names that reference a physical server name. For example, rename the `.sapenv_physicalserver.csh` file as `.sapenv.csh`.

5. As user `sapsidadm`, add an entry such as the following example for the parameter `SAPLOCALHOST`.

Add this entry to the `SAPSID_Service-StringSystem-Number_logical-hostname` profile file under the `/sapmnt/SAPSID/profile` directory.

```
SAPLOCALHOST=ci-logical-hostname
```

This entry enables the external application server to locate the central instance by using the network resource (logical hostname or shared address).

6. Become user `orasapsid`, and then perform the following tasks.
 - In the `oraSAPSID` home directory, modify all of the file names that reference a physical server name.
 - In the `oraSAPSID` home directory, modify all of the file contents—except log file contents—that reference a physical server name.
 - In the SAP profile directory, modify all of the file names that reference a physical server name. For example, change any *physicalserver* reference in the startup and shutdown scripts to a logical hostname reference.

Where to Go From Here

Go to [“How to Verify SAP and the Database Installation with Central Instance”](#) on page 45.

▼ How to Verify SAP and the Database Installation with Central Instance

To test starting and stopping the SAP central instance on all of the potential nodes on which the central instance can run, perform the following steps.

1. Create the failover resource group to hold the network and central instance resources.

```
# scrgadm -a -g sap-ci-resource-group
```

Note – Use the `-h` option to the `scrgadm(1M)` command to select the set of nodes on which the SAP central instance can run.

```
# scrgadm -a -g sap-ci-resource-group [-h nodelist]
```

2. Verify that you have added to your name service database all of the network resources that you use.
3. Add a network resource (logical hostname) to the failover resource group.

```
# scrgadm -a -L -g sap-ci-resource-group -l ci-logical-hostname [-n netiflist]
```

4. Enable the resource group.

Run the `scswitch(1M)` command to move the resource group into a managed state and bring the resource group online.

```
# scswitch -Z -g sap-ci-resource-group
```

5. Log in to the cluster member that hosts the central instance resource group.
6. Start the database.
7. Start the central instance.
8. Start the SAP GUI to verify that SAP initializes correctly.
The default dispatcher port is 3200.
9. Stop the central instance.
10. Run the `scswitch` command.

In the following example, the replaceable term *sap-ci-resource-group* represents the resource group that contains the network resource (logical hostname or shared address) for the central instance resource. Switch this resource group to another cluster member that can host the central instance.

```
# scswitch -z -h node -g sap-ci-resource-group
```

11. Repeat [Step 5](#) through [Step 9](#) until you verify startup and shutdown of the central instance on each cluster node that can host the central instance.

Where to Go From Here

Go to [“How to Verify SAP and the Database Installation with Failover Application Server”](#) on page 48.

▼ How to Verify SAP and the Database Installation with Failover Application Server

To test starting and stopping the failover application server, perform the following steps procedure on all of the potential nodes on which the failover application server can run.

1. **Create the failover resource group to hold the network and application server resources.**

```
# scrgadm -a -g sap-as-fo-resource-group
```

Note – Use the `-h` option to the `scrgadm` command to select the set of nodes on which the SAP application server can run.

```
# scrgadm -a -g sap-as-fo-resource-group [-h nodelist]
```

2. **Verify that you added to your name service database all of the network resources that you use.**
3. **Add a network resource (logical hostname) to the failover resource group.**

```
# scrgadm -a -L -g sap-as-fo-resource-group -l as-fo-logical-hostname [-n netiflist]
```

4. **Enable the resource group.**

Run the `scswitch(1M)` command to move the resource group into a managed state and bring the resource group online.

```
# scswitch -Z -g sap-as-of-resource-group
```

5. **Log in to the cluster member that hosts the application server resource group.**
6. **Start the application server.**
7. **Start the SAP GUI to verify that the SAP application server initializes correctly.**
8. **Stop the application server.**

9. Run the `scswitch` command.

In the following example, the term *sap-as-resource-group* represents the resource group that contains the network resource (logical hostname or shared address) for the application server resource. Switch this resource group to another cluster member that can host the application server.

```
# scswitch -z -h node -g sap-as-fo-resource-group
```

- 10. Repeat [Step 5](#) through [Step 7](#) until you verify startup and shutdown of the application server on each cluster node that can host the application server.**

Where to Go From Here

Go to [“Configuring Sun Cluster HA for DBMS” on page 50](#) for information on how to configure DBMS, an SAP service, so that it is highly available.

Configuring Sun Cluster HA for DBMS

SAP supports various databases. See the appropriate chapter of this book for details on how to configure the resource type, resource group, and resource for your highly available database. For example, see “Installing and Configuring Sun Cluster HA for Oracle” in the *Sun Cluster 3.0 12/01 Data Services Installation and Configuration Guide* for more information if you plan to use Oracle with SAP.

Additionally, see the appropriate chapter of this book and the appropriate chapter of your database installation book for details on other resource types to configure with your database. This book includes details on how to configure other resource types for Oracle databases. For instance, set up the `SUNW.HASStorage` resource type if you use Oracle. See the procedure “How to Configure SUNW.HASStorage Resource Type” in the *Sun Cluster 3.0 12/01 Data Services Installation and Configuration Guide* for more information.

Where to Go From Here

Go to [“Registering and Configuring Sun Cluster HA for SAP”](#) on page 51.

Registering and Configuring Sun Cluster HA for SAP

Use the procedures in this section to perform the following tasks.

- Register and configure Sun Cluster HA for SAP central instance.
- Register and configure Sun Cluster HA for SAP with a failover application server.
- Register and configure Sun Cluster HA for SAP with a scalable application server.

▼ How to Register and Configure Sun Cluster HA for SAP Central Instance

To register and configure Sun Cluster HA for SAP central instance, perform the following steps.

1. **Become superuser on one of the nodes in the cluster that hosts the central instance.**
2. **Register the resource type for the central instance.**

```
# scrgadm -a -t SUNW.sap_ci | SUNW.sap_ci_v2
```

3. **Create SAP application server resources in this failover resource group.**

```
# scrgadm -a -j sap-ci-resource -g sap-ci-resource-group \  
-t SUNW.sap_ci | SUNW.sap_ci_v2 \  
-x SAPSID=SAPSID \  
-x Ci_startup_script=ci-startup-script \  
-x Ci_shutdown_script=ci-shutdown-script
```

See “Configuring Sun Cluster HA for SAP Extension Properties” in the *Sun Cluster 3.0 12/01 Data Services Installation and Configuration Guide* for a list of extension properties.

4. **Enable the failover resource group that now includes the SAP central instance resource.**

```
# scswitch -Z -g sap-ci-resource-group
```

Where to Go From Here

Go to [“How to Register and Configure Sun Cluster HA for SAP with Failover Application Server”](#) on page 52 or [“How to Register and Configure Sun Cluster HA for SAP with Scalable Application Server”](#) on page 53.

▼ How to Register and Configure Sun Cluster HA for SAP with Failover Application Server

To register and configure Sun Cluster HA for SAP with a failover application server, perform the following steps.

1. **Become superuser on one of the nodes in the cluster that hosts the application server.**
2. **Register the resource type for the failover application server.**

```
# scrgadm -a -t SUNW.sap_as | SUNW.sap_as_v2
```

3. **Create SAP application server resources in this failover resource group.**

```
# scrgadm -a -j sap-as-resource -g sap-as-fo-resource-group \  
-t SUNW.sap_as | SUNW.sap_as_v2 \  
-x SAPSID=SAPSID \  
-x As_instance_id=as-instance-id \  
-x As_startup_script=as-startup-script \  
-x As_shutdown_script=as-shutdown-script
```

See [“Configuring Sun Cluster HA for SAP Extension Properties”](#) on page 58 for a list of extension properties.

4. **Enable the failover resource group that now includes the SAP application server resource.**

```
# scswitch -Z -g sap-as-fo-resource-group
```

Where to Go From Here

Go to [“How to Verify the Installation and Configuration of Sun Cluster HA for SAP, Central Instance, and DBMS”](#) on page 55.

▼ How to Register and Configure Sun Cluster HA for SAP with Scalable Application Server

To register and configure Sun Cluster HA for SAP with a scalable application server, perform the following steps.

1. Become superuser on one of the nodes in the cluster that hosts the application server.
2. Create a scalable resource group for the application server.

```
# scrgadm -a -g sap-as-sa-appinstanceid-resource-group \  
-y Maximum primaries=value \  
-y Desired primaries=value
```

Note – You will not use the shared address to perform load balancing of the application server. The SAP Message Server will be used to perform load balancing of the application server.

Note – If you are using the `SUNW.RGoffload` resource type to offload an application server within this scalable application server resource group, then set `Desired primaries=0`. See “Freeing Node Resources by Offloading Non-critical Resource Groups” in the *Sun Cluster 3.0 12/01 Data Services Installation and Configuration Guide* for more information about using the `SUNW.RGoffload` resource type.

3. Register the resource type for the scalable application server.

```
# scrgadm -a -t SUNW.sap_as_v2
```

4. Create SAP application server resources in this scalable resource group.

```
# scrgadm -a -j sap-as-resource -g sap-as-sa-appinstanceid-resource-group \  
-t SUNW.sap_as_v2 \  
-x SAPSID=SAPSID \  
-x As_instance_id=as-instance-id \  
-x As_startup_script=as-startup-script \  
-x As_shutdown_script=as-shutdown-script
```

See [“Configuring Sun Cluster HA for SAP Extension Properties” on page 58](#) for a list of extension properties.

5. Enable the scalable resource group that now includes the SAP application server resource.

- If you *do not* use the `SUNW.RGoffload` resource type with this application server, use the following command.

```
# scswitch -Z -g sap-as-sa-appinstanceid-resource-group
```

- If you use the `SUNW.RGoffload` resource type with this application server, use the following command.

```
# scswitch -z -h node1, node2 -g sap-as-sa-appinstanceid-resource-group
```

Note – If you use the `SUNW.RGoffload` resource type with this application server, you must specify which node you want to bring the resource online by using the `-z` option instead of the `-j` option.

Where to Go From Here

Go to [“How to Verify the Installation and Configuration of Sun Cluster HA for SAP, Central Instance, and DBMS” on page 55](#).

Verifying Sun Cluster HA for SAP Installation and Configuration

Use the procedures in this section to perform the following tasks.

- Verify the installation and configuration of Sun Cluster HA for SAP, the central instance, and DBMS.
- Verify the installation and configuration of Sun Cluster HA for SAP with a failover application server.
- Verify the installation and configuration of Sun Cluster HA for SAP with a scalable application server.

▼ How to Verify the Installation and Configuration of Sun Cluster HA for SAP, Central Instance, and DBMS

To verify both the Sun Cluster HA for SAP installation with the central instance and the DBMS installation and configuration, perform the following steps.

1. **Log in to the node that hosts the resource group that contains the SAP central instance resource.**
2. **Become** *user sapsidadm*.
3. **Start the SAP GUI to check that Sun Cluster HA for SAP is functioning correctly.**
4. **Use the central instance** `stopsap` **script to shut down the SAP central instance.**
The Sun Cluster software should restart the central instance because the Sun Cluster software controls the SAP software.
5. **Switch the SAP resource group to another cluster member.**

```
# scswitch -z -h node2 -g sap-ci-resource-group
```

6. **Verify that the SAP central instance starts on this node.**
7. **Repeat [Step 1](#) through [Step 6](#) until you have tested all of the potential nodes on which the SAP central instance can run.**

Where to Go From Here

Go to “How to Verify the Installation and Configuration of Sun Cluster HA for SAP with Failover Application Server” on page 56 or “How to Verify the Installation and Configuration of Sun Cluster HA for SAP with Scalable Application Server” on page 57.

▼ How to Verify the Installation and Configuration of Sun Cluster HA for SAP with Failover Application Server

To verify the Sun Cluster HA for SAP installation and configuration with a failover application server, perform the following steps.

1. **Log in to the node that currently hosts the resource group that contains the SAP application server resource.**
2. **Become *user sapsidadm*.**
3. **Start the SAP GUI to check that the application server is functioning correctly.**
4. **Use the application server `stopsap` script to shut down the SAP application server on the node you identified in [Step 1](#).**

The Sun Cluster software should restart the application server because the Sun Cluster software controls the SAP software.

5. **Switch the resource group that contains the SAP application server resource to another cluster member.**

```
# scswitch -z -h node2 -g sap-as-resource-group
```

6. **Verify that the SAP application server starts on the node you identified in [Step 5](#).**
7. **Repeat [Step 1](#) through [Step 6](#) until you have tested all of the potential nodes on which the SAP application server can run.**

▼ How to Verify the Installation and Configuration of Sun Cluster HA for SAP with Scalable Application Server

To verify the Sun Cluster HA for SAP installation and configuration with a scalable application server, perform the following steps.

1. Log on to one of the nodes that runs the application server.
2. Become *user sapsidadm*.
3. Start the SAP GUI to check that the application server is functioning correctly.
4. Use the application server `stopsap` script to shut down the SAP application server on the node you identified in [Step 1](#).

The Sun Cluster software should restart the application server because the Sun Cluster software controls the SAP software.

5. Repeat [Step 1](#) through [Step 4](#) until you have tested all of the potential nodes on which the SAP application server can run.

Configuring Sun Cluster HA for SAP Extension Properties

Use the Sun Cluster HA for SAP extension properties in this section to create the central instance resources and application server resources. Typically, you use the command line `scrgadm -x parameter=value` to configure the extension properties when you create the central instance or application resource. You can also use the procedures described in Chapter 13 of the *Sun Cluster 3.0 12/01 Data Services Installation and Configuration Guide* to configure them later. See Appendix A of the *Sun Cluster 3.0 12/01 Data Services Installation and Configuration Guide* for details on all of the Sun Cluster properties.

See the `r_properties(5)` and the `rg_properties(5)` man pages for details on all of the Sun Cluster extension properties.

[TABLE B-5](#) describes SAP extension properties that you can set for the central instance. You can update some extension properties dynamically. You can update others, however, only when you create or disable the SAP resource. The Tunable entries indicate when you can update each property. [TABLE B-6](#) describes the extension properties that you can set for the application servers.

TABLE B-5 Sun Cluster HA for SAP Extension Properties for the Central Instance (Sheet 1 of 4)

Property Category	Property Name	Description
SAP Configuration	SAPSID	SAP system name or <i>SAPSID</i> . Default: None Tunable: When disabled
	Ci_instance_id	Two-digit SAP system number. Default: 00 Tunable: When disabled
	Ci_services_string	String of central instance services. Default: DVEBMGS Tunable: When disabled

TABLE B-5 Sun Cluster HA for SAP Extension Properties for the Central Instance (Sheet 2 of 4)

Property Category	Property Name	Description
Starting SAP	ci_start_retry_interval	The interval in seconds to wait between attempting to connect to the database before starting the central instance. Default: 30 Tunable: When disabled
	ci_startup_script	Name of the SAP startup script for this instance in your <i>SIDadm</i> home directory. Default: None Tunable: When disabled
Stopping SAP	stop_sap_pct	Percentage of stop-timeout variables that are used to stop SAP processes. The SAP shutdown script is used to stop processes before calling Process Monitor Facility (PMF) to terminate and then kill the processes. Default: 95 Tunable: When disabled
	ci_shutdown_script	Name of the SAP shutdown script for this instance in your <i>SIDadm</i> home directory. Default: None Tunable: When disabled

TABLE B-5 Sun Cluster HA for SAP Extension Properties for the Central Instance (Sheet 3 of 4)

Property Category	Property Name	Description
Probe	Message_server_name	The name of the SAP Message Server. Default: <code>sapms SAPSID</code> Tunable: When disabled
	Lgtst_ms_with_logicalhostname	How to check the SAP Message Server with the SAP <code>lgtst</code> utility. The <code>lgtst</code> utility requires a hostname (IP address) as the location for the SAP Message Server. This hostname can be either a Sun Cluster logical hostname or a localhost (loopback) name. If you set this resource property to <code>TRUE</code> , use a logical hostname. Otherwise, use a localhost name. Default: <code>TRUE</code> Tunable: Any time
	Check_ms_retry	Maximum number of times the SAP Message Server check fails before a total failure is reported and the Resource Group Manager (RGM) starts. Default: 2 Tunable: When disabled
	Probe_timeout	Time-out value in seconds for the probes. Default: 60 Tunable: Any time
	Monitor_retry_count	Number of PMF restarts that are allowed for the fault monitor. Default: 4 Tunable: Any time
	Monitor_retry_interval	Time interval in minutes for fault monitor restarts. Default: 2 Tunable: Any time

TABLE B-5 Sun Cluster HA for SAP Extension Properties for the Central Instance (Sheet 4 of 4)

Property Category	Property Name	Description
Development System	Shutdown_dev	Whether the RGM should shut down the development system before starting up the central instance. Default: FALSE Tunable: When disabled
	Dev_sapsid	SAP System Name for the development system (if you set Shutdown_dev to TRUE, Sun Cluster HA for SAP requires this property). Default: None Tunable: When disabled
	Dev_shutdown_script	Script that is used to shut down the development system. If you set Shutdown_dev to TRUE, Sun Cluster HA for SAP requires this property. Default: None Tunable: When disabled
	Dev_stop_pct	Percentage of startup timeouts Sun Cluster HA for SAP uses to shut down the development system before starting the central instance. Default: 20 Tunable: When disabled

TABLE B-6 Sun Cluster HA for SAP Extension Properties for the Application Servers (Sheet 1 of 2)

Property Category	Property Name	Description
SAP Configuration	SAPSID	SAP system name or <i>SAPSID</i> for the application server. Default: None Tunable: When disabled
	As_instance_id	Two-digit SAP system number for the application server. Default: None Tunable: When disabled
	As_services_string	String of application server services. Default: D Tunable: When disabled

TABLE B-6 Sun Cluster HA for SAP Extension Properties for the Application Servers
(Sheet 2 of 2)

Property Category	Property Name	Description
Starting SAP	As_db_retry_interval	The interval in seconds to wait between attempting to connect to the database and starting the application server. Default: 30 Tunable: When disabled
	As_startup_script	Name of the SAP startup script for the application server. Default: None Tunable: When disabled
Stopping SAP	Stop_sap_pct	Percentage of stop-timeout variables that are used to stop SAP processes. The SAP shutdown script is used to stop processes before calling Process Monitor Facility (PMF) to terminate and then kill the processes. Default: 95 Tunable: When disabled
	As_shutdown_script	Name of the SAP shutdown script for the application server. Default: None Tunable: When disabled
Probe	Probe_timeout	Time-out value in seconds for the probes. Default: 60 Tunable: Any time
	Monitor_retry_count	Number of PMF restarts that the probe allows for the fault monitor. Default: 4 Tunable: Any time
	Monitor_retry_interval	Time interval in minutes for fault monitor restarts. Default: 2 Tunable: Any time

Setting Up a Lock File

Use the procedure in this section to perform the following tasks.

- Set up a lock file for the central instance.
- Set up a lock file for a scalable application server.

Set up a lock file to prevent multiple startups of the SAP instance when the instance is already active on one node. Multiple startups of the same instance crash each other. Furthermore, the crash prevents SAP shutdown scripts from performing a clean shutdown of the instances, which might cause data corruption.

If you set up a lock file, when you start the SAP instance the SAP software locks the file `startup_lockfile`. If you start up the same instance outside of the Sun Cluster environment and then try to bring up SAP under the Sun Cluster environment, the Sun Cluster HA for SAP data service will attempt to start up the same instance. However, because of the file-locking mechanism, this attempt will fail. The data service will log appropriate error messages in `syslog`.

The only difference between the lock file for the central instance or the failover application server and the lock file for a scalable application server is that the lock file for scalable application server resides on the local file system and the lock file for the central instance or the failover application server resides on a cluster file system.

▼ How to Set Up a Lock File for Central Instance

To set up a central instance lock file, perform the following steps.

1. **Install the latest patch for the `sapstart` executable, which enables Sun Cluster HA for SAP users to configure a lock file.**
2. **Set up the central instance lock file on a *cluster* file system.**

3. Edit the profile `SC3_DVEBMGS00` to add the new SAP parameter, `sapstart/lockfile`, for the application server.

```
sapstart/lockfile =/usr/sap/SC3/DVEBMGS00/work/startup_lockfile
```

<code>sapstart/lockfile</code>	New parameter name.
<code>/usr/sap/local/SC3/DVEBMGS00/work</code>	Work directory for the application server.
<code>startup_lockfile</code>	Lock file name that Sun Cluster HA for SAP uses.

SAP creates the lock file.

Note – You must locate the lock file path on a cluster file system. If you locate the lock file path locally on the nodes, startups of the same instance from multiple nodes cannot be prevented.

▼ How to Set Up a Lock File for Scalable Application Server

To set up a lock file for a scalable application server, perform the following steps.

1. Install the latest patch for the `sapstart` executable, which enables Sun Cluster HA for SAP users to configure a lock file.
2. Set up the application server lock file on the *local* file system.

3. Edit the profile `SC3_instance-id` to add the new SAP parameter, `sapstart/lockfile`, for the application server.

```
sapstart/lockfile =/usr/sap/local/SC3/Dinstance-id/work/startup_lockfile
```

<code>sapstart/lockfile</code>	New parameter name.
<code>/usr/sap/SC3/Dinstance-id/work</code>	Work directory for the central instance.
<code>startup_lockfile</code>	Lock file name that Sun Cluster HA for SAP uses.

SAP creates the lock file.

Note – The lock file will reside on the local file system. The lock file does not prevent multiple startups from *other* nodes, but the lock file does prevent multiple startups from the *same* node.

Sun Cluster HA for SAP Fault Monitor

The Sun Cluster HA for SAP fault monitor checks SAP process and database health. SAP process health impacts SAP resources' failure history. SAP resources' failure history in turn drives the fault monitor's actions, which include no action, restart, or failover.

In contrast to SAP process health, the health of the database SAP uses has no impact on SAP resources' failure history. Database health does, however, trigger the SAP fault monitor to log any `syslog` messages and to set the status accordingly for the SAP resource that uses the database.

Sun Cluster HA for SAP Fault Probes for Central Instance

For the central instance, the fault probe executes the following steps.

1. Retrieves the process IDs for the SAP Message Server and the dispatcher
2. Loops infinitely (sleeps for `Thorough_probe_interval`)
3. Checks the health of the SAP resources
 - a. **Abnormal exit** – If the Process Monitor Facility (PMF) detects that the SAP process tree has failed, the fault monitor treats this problem as a complete failure. The fault monitor restarts or fails over the SAP resource to another node based on the resources' failure history.
 - b. **Health check of the SAP resources through probe** – The probe uses the `ps(1)` command to check the SAP Message Server and main dispatcher processes. If any of the SAP Message Server or main dispatcher processes are missing from the system's active processes list, the fault monitor treats this problem as a complete failure.

If you configure the parameter `Check_ms_retry` to have a value greater than zero, the probe checks the SAP Message Server connection. If you have set the extension property `Lgtst_ms_with_logicalhostname` to its default value `TRUE`, the probe completes the SAP Message Server connection test with the utility `lgtst`. The probe uses the logical hostname interface that is specified in the SAP resource group to call the SAP-supplied utility `lgtst`. If you set the extension property `Lgtst_ms_with_logicalhostname` to a value other than `TRUE`, the probe calls `lgtst` with the node's localhost name (loopback interface).

If the `lgtst` utility call fails, the SAP Message Server connection is not functioning. In this situation, the fault monitor considers the problem to be a partial failure and does not trigger an SAP restart or a failover immediately. The fault monitor counts two partial failures as a complete failure if the following conditions occur.

- i. You configure the extension property `Check_ms_retry` to be 2.
- ii. The fault monitor accumulates two partial failures that happen within the retry interval that the resource property `Retry_interval` sets.

A complete failure triggers either a local restart or a failover, based on the resource's failure history.

- c. **Database connection status through probe** – The probe calls the SAP-supplied utility `R3trans` to check the status of the database connection. Sun Cluster HA for SAP fault probes verify that SAP can connect to the database. Sun Cluster HA for SAP depends, however, on the highly available database fault probes to determine the health of the database. If the database connection status check fails, the fault monitor logs the message, `Database might be down`, to `syslog`. The fault monitor then sets the status of the SAP resource to `DEGRADED`. If the probe checks the status of the database again and the connection is reestablished, the fault monitor logs the message, `Database is up`, to `syslog` and sets the status of the SAP resource to `OK`.

4. Evaluates the failure history

Based on the failure history, the fault monitor completes one of the following actions.

- no action
- local restart
- failover

Sun Cluster HA for SAP Fault Probes for Application Server

For the application server, the fault probe executes the following steps.

1. Retrieves the process ID for the main dispatcher
2. Loops infinitely (sleeps for `Thorough_probe_interval`)

3. Checks the health of the SAP resources

- a. **Abnormal exit** – If the Process Monitor Facility (PMF) detects that the SAP process tree has failed, the fault monitor treats this problem as a complete failure. The fault monitor restarts or fails over the SAP resource to another node, based on the resources' failure history.
- b. **Health check of the SAP resources through probe** – The probe uses the `ps(1)` command to check the SAP Message Server and main dispatcher processes. If the SAP main dispatcher process is missing from the system's active processes list, the fault monitor treats the problem as a complete failure.
- c. **Database connection status through probe** – The probe calls the SAP-supplied utility `R3trans` to check the status of the database connection. Sun Cluster HA for SAP fault probes verify that SAP can connect to the database. Sun Cluster HA for SAP depends, however, on the highly available database fault probes to determine the health of the database. If the database connection status check fails, the fault monitor logs the message, `Database might be down`, to `syslog` and sets the status of the SAP resource to `DEGRADED`. If the probe checks the status of the database again and the connection is reestablished, the fault monitor logs the message, `Database is up`, to `syslog`. The fault monitor then sets the status of the SAP resource to `OK`.

4. Evaluates the failure history

Based on the failure history, the fault monitor completes one of the following actions.

- no action
- local restart
- failover

If the application server resource is a failover resource, the fault monitor fails over the application server.

If the application server resource is a scalable resource, after the number of local restarts are exhausted, RGM will bring up the application server on a different node if there is another node available in the cluster.

Installing and Maintaining a Sun StorEdge 3900 Series System

This chapter contains the procedures for installing, configuring, and maintaining Sun StorEdge 3910 and Sun StorEdge 3960 systems.

This chapter contains the following procedures:

- [“How to Install StorEdge 3900 Series Systems” on page 70](#)
- [“How to Configure StorEdge 3900 Series Logical Volumes” on page 73](#)
- [“How to Add StorEdge 3900 Series Systems to a Running Cluster” on page 77](#)
- [“How to Remove StorEdge 3900 Series Systems From a Running Cluster” on page 85](#)
- [“How to Replace a Node-to-Switch Component in a Running Cluster” on page 91](#)
- [“How to Replace an FC Switch or Array-to-Switch Component in a Running Cluster” on page 88](#)
- [“How to Remove StorEdge T3/T3+ Arrays From a Running Cluster” on page 92](#)
- [“How to Upgrade StorEdge T3+ Array Firmware” on page 95](#)
- [“How to Replace a Failed Disk Drive in a Running Cluster” on page 98](#)
- [“How to Replace a Node’s Host Adapter in a Running Cluster” on page 99](#)

The StorEdge 3900 Series configuration utilities can be run from a menu-driven interface or a command line interface (this chapter describes the menu-driven interface). For detailed information about Sun StorEdge 3900 Series architecture, features, and configuration utilities, see the *Sun StorEdge 3900 and 6900 Series Reference Manual* and the *Sun StorEdge 3900 and 6900 Series Cabinet Installation and Service Manual*.

Installing StorEdge 3900 Series Systems

Note – This section contains the procedure for an *initial installation* of StorEdge 3900 Series systems in a new Sun Cluster that is not running. If you are adding systems to an existing cluster, use the procedure in [“How to Add StorEdge 3900 Series Systems to a Running Cluster”](#) on page 77.

▼ How to Install StorEdge 3900 Series Systems

Perform the steps in this procedure in conjunction with the procedures in the *Sun Cluster 3.0 12/01 Software Installation Guide* and your server hardware manual.

1. Install host adapters in the nodes that will be connected to the system.

For the procedure on installing host adapters, see the documentation that shipped with your host adapters and nodes.

2. Unpack, place, and level the system cabinet.

For instructions, see the *Sun StorEdge 3900 and 6900 Series Cabinet Installation and Service Manual*.

3. Install cables in the following order.

a. Install the system power cord.

b. Install the system grounding strap.

c. Install the cables from the FC switches to the cluster nodes.

d. Install the Ethernet cable to the local area network (LAN).

For instructions on cabling, see the *Sun StorEdge 3900 and 6900 Series Cabinet Installation and Service Manual*.

4. Power on the system.

For instructions, see the *Sun StorEdge 3900 and 6900 Series Cabinet Installation and Service Manual*.

5. Set the host name, IP address, date, and timezone for the system’s storage service processor.

For detailed instructions, see the initial field installation instructions in the *Sun StorEdge 3900 and 6900 Series Reference Manual*.

6. Remove the preconfigured, default hard zoning from the system's FC switches.

Note – To configure the StorEdge 3900 Series system for use with Sun Cluster host-based mirroring, the default hard zones *must* be removed from the system's FC switches. See the *SANbox-8/16 Switch Management User's Manual* for instructions on using the installed SANSurfer interface for removing the preconfigured hard zones from all Sun StorEdge Network FC Switch-8 and Switch-16 switches.

7. If necessary, upgrade StorEdge T3+ array controller firmware.

Telnet to the storage service processor and install the required controller firmware. For instruction on upgrading StorEdge T3+ firmware, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

See the *Sun Cluster 3.0 12/01 Release Notes* for information about accessing Sun's EarlyNotifier web pages, which list information about any required patches or firmware levels that are available for download. For the procedure on applying any host adapter firmware patch, see the firmware patch README file.

8. (Optional) Configure the arrays with the desired logical volumes.

To configure logical volumes, use the StorEdge 3900 Series' menu-driven or command-line interface, as described in the *Sun StorEdge 3900 and 6900 Series Reference Manual*.

9. Install the Solaris operating environment to the cluster nodes and apply the required Solaris patches for Sun Cluster software and StorEdge T3+ array support.

For the procedure on installing the Solaris operating environment, see the *Sun Cluster 3.0 12/01 Software Installation Guide*.

See the *Sun Cluster 3.0 12/01 Release Notes* for information about accessing Sun's EarlyNotifier web pages, which list information about any required patches or firmware levels that are available for download. For the procedure on applying any host adapter firmware patch, see the firmware patch README file.

10. Install any required patches or software for Sun StorEdge Traffic Manager software support to the cluster nodes from the Sun Download Center Web site, <http://www.sun.com/storage/san/>.

For instructions on installing the software, see the information on the web site.

11. Activate the Sun StorEdge Traffic Manager software functionality in the software you installed to the cluster nodes in [Step 10](#).

To activate the Sun StorEdge Traffic Manager software functionality, manually edit the `/kernel/drv/scsi_vhci.conf` file that is installed on the node to change the `mpxio-disable` parameter to no:

```
mpxio-disable="no"
```

12. Perform a reconfiguration boot on all nodes to create the new Solaris device files and links.

```
{0} ok boot -r
```

13. On all nodes, update the `/devices` and `/dev` entries:

```
# devfsadm -C
```

14. On all nodes, use the `luxadm display` command to confirm that all arrays you installed are now visible.

```
# luxadm display
```

Where to Go From Here

To continue with Sun Cluster software installation tasks, see the *Sun Cluster 3.0 12/01 Software Installation Guide*.

Configuring StorEdge 3900 Series Systems in a Running Cluster

To configure logical volumes on the StorEdge T3+ arrays in your system, use the StorEdge 3900 Series' menu-driven or command-line interface, as described in the *Sun StorEdge 3900 and 6900 Series Reference Manual*.

▼ How to Configure StorEdge 3900 Series Logical Volumes

A summary of the steps involved with configuring the StorEdge T3+ arrays in your StorEdge 3900 system is listed below.

- If you are using the menu-driven interface, perform the following steps.

1. Start the Configuration Utilities using the `runsecfg` command.
2. Select `T3+ Configuration Utility` from the main menu.
3. Enter your StorEdge T3+ array password when you are prompted.
4. Step through the submenus to configure the StorEdge T3+ arrays. See the *Sun StorEdge 3900 and 6900 Series Reference Manual* for more information about the submenus.

- If you are using the command-line interface, do the following steps.

1. If prompted, enter your StorEdge T3+ array password.
2. Use the `modifyt3config` command to configure the arrays in your system. See the *Sun StorEdge 3900 and 6900 Series Reference Manual* and the `man` page for this command for more information.

Maintaining StorEdge 3900 Series Systems

This section contains the procedures for maintaining StorEdge 3900 Series systems. [TABLE C-1](#) lists these procedures. This section does not include procedures for adding or removing disk drives because the StorEdge T3+ arrays in your StorEdge 3900 Series system operate only when fully configured with disk drives.



Caution – If you remove any field replaceable unit (FRU) from the StorEdge T3+ arrays for an extended period of time, thermal complications might result. To prevent these complications, the StorEdge T3+ array is designed so that an orderly shutdown occurs when you remove a component for longer than 30 minutes. Therefore, a replacement part must be immediately available before you start an FRU replacement procedure. You must replace an FRU within 30 minutes or the StorEdge T3+ array, and all attached StorEdge T3+ arrays, will shut down and power off.

TABLE C-1 Task Map: Maintaining a StorEdge 3900 Series System (Sheet 1 of 3)

Task	For Instructions, Go To...
Add a StorEdge 3900 Series system.	"How to Add StorEdge 3900 Series Systems to a Running Cluster" on page 77
Remove a StorEdge 3900 Series system.	"How to Remove StorEdge 3900 Series Systems From a Running Cluster" on page 85
Replace a StorEdge network FC switch-8 or switch-16.	"How to Replace an FC Switch or Array-to-Switch Component in a Running Cluster" on page 88
Replace an array-to-switch fiber optic cable.	"How to Replace an FC Switch or Array-to-Switch Component in a Running Cluster" on page 88
Replace a GBIC on an FC switch, connecting to an array.	"How to Replace an FC Switch or Array-to-Switch Component in a Running Cluster" on page 88
Replace a StorEdge network FC switch-8 or switch-16 power cord.	"How to Replace an FC Switch or Array-to-Switch Component in a Running Cluster" on page 88

TABLE C-1 Task Map: Maintaining a StorEdge 3900 Series System (Sheet 2 of 3)

Task	For Instructions, Go To...
Replace interconnect cables.	"How to Replace an FC Switch or Array-to-Switch Component in a Running Cluster" on page 88
Replace a StorEdge T3+ array controller.	"How to Replace an FC Switch or Array-to-Switch Component in a Running Cluster" on page 88
Replace a node-to-switch fiber optic cable.	"How to Replace a Node-to-Switch Component in a Running Cluster" on page 91
Replace a gigabit interface converter (GBIC) on a node's host adapter.	"How to Replace a Node-to-Switch Component in a Running Cluster" on page 91
Replace a GBIC on an FC switch, connecting to a node.	"How to Replace a Node-to-Switch Component in a Running Cluster" on page 91
Remove a T3+ array partner group from the system.	"How to Remove StorEdge T3/T3+ Arrays From a Running Cluster" on page 92
Upgrade StorEdge T3+ array firmware.	"How to Upgrade StorEdge T3+ Array Firmware" on page 95
Replace a disk drive in an array.	"How to Replace a Failed Disk Drive in a Running Cluster" on page 98
Replace a host adapter in a node.	"How to Replace a Node's Host Adapter in a Running Cluster" on page 99
Replace a StorEdge T3+ Power and Cooling Unit (PCU).	<i>Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual</i>
Follow the same procedure used in a non-cluster environment.	
Replace a StorEdge T3+ unit interconnect card (UIC).	<i>Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual</i>
Follow the same procedure used in a non-cluster environment.	
Replace a StorEdge T3+ array power cable.	<i>Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual</i>
Follow the same procedure used in a non-cluster environment.	

TABLE C-1 Task Map: Maintaining a StorEdge 3900 Series System (Sheet 3 of 3)

Task	For Instructions, Go To...
Replace an Ethernet hub. Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge 3900 and 6900 Series Reference Manual</i>
Replace a StorEdge T3+ Ethernet cable. Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual</i>
Replace a storage service processor. Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge 3900 and 6900 Series Reference Manual</i>

▼ How to Add StorEdge 3900 Series Systems to a Running Cluster

Note – Use this procedure to add new StorEdge 3900 Series systems to a running cluster. To install systems to a new Sun Cluster that is not running, use the procedure in [“How to Install StorEdge 3900 Series Systems” on page 70](#).

This procedure defines “Node A” as the node you begin working with, and “Node B” as the second node.

1. Unpack, place, and level the system cabinet.

For instructions, see the *Sun StorEdge 3900 and 6900 Series Cabinet Installation and Service Manual*.

2. Install cables in the following order.

a. Install the system power cord.

b. Install the system grounding strap.

c. Install the cables from the FC switches to the cluster nodes.

d. Install the Ethernet cable to the LAN.

For instructions on cabling, see the *Sun StorEdge 3900 and 6900 Series Cabinet Installation and Service Manual*.

3. Power on the new system.

Note – The StorEdge T3+ arrays in your system might take several minutes to boot.

For instructions, see the *Sun StorEdge 3900 and 6900 Series Cabinet Installation and Service Manual*.

4. Set the host name, IP address, date, and timezone for the system’s Storage Service Processor.

For detailed instructions, see the initial field installation instructions in the *Sun StorEdge 3900 and 6900 Series Reference Manual*.

5. Remove the preconfigured, default hard zoning from the new system's FC switches.

Note – To configure the StorEdge 3900 Series system for use with Sun Cluster host-based mirroring, the default hard zones *must* be removed from the system's FC switches. See the *SANbox-8/16 Switch Management User's Manual* for instructions on using the installed SANsurfer interface for removing the preconfigured hard zones from all Sun StorEdge Network FC Switch-8 and Switch-16 switches.

6. Install any required StorEdge T3+ array controller firmware upgrades:

Telnet to the storage service processor and install the required controller firmware. For instruction on upgrading StorEdge T3+ firmware, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

See the *Sun Cluster 3.0 12/01 Release Notes* for information about accessing Sun's EarlyNotifier web pages, which list information about any required patches or firmware levels that are available for download. For the procedure on applying any host adapter firmware patch, see the firmware patch README file.

7. (Optional) Configure the new arrays with the desired logical volumes.

For the procedure on creating and initializing a logical volume, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*. For the procedure on mounting a logical volume, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

8. Determine the resource groups and device groups running on all nodes.

Record this information because you will use it in [Step 47](#) of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

9. Move all resource groups and device groups off Node A.

```
# scswitch -s -h nodename
```

10. Do you need to install host adapters in Node A?

- If not, go to [Step 16](#).
- If you do need to install host adapters to Node A, continue with [Step 11](#).

11. Is the host adapter you are installing the first host adapter on Node A?

- If not, go to [Step 13](#).

- If it is the first host adapter, use the `pkginfo` command as shown below to determine whether the required support packages for the host adapter are already installed on this node. The following packages are required.

```
# pkginfo | egrep Wlux
system    SUNWluxd   Sun Enterprise Network Array sf Device Driver
system    SUNWluxdx  Sun Enterprise Network Array sf Device Driver (64-bit)
system    SUNWluxl   Sun Enterprise Network Array social Device Driver
system    SUNWluxlx  Sun Enterprise Network Array social Device Driver (64-bit)
system    SUNWluxop  Sun Enterprise Network Array firmware and utilities
system    SUNWluxox  Sun Enterprise Network Array libraries (64-bit)
```

12. Are the required support packages already installed?

- If they are already installed, go to [Step 13](#).
- If not, install the required support packages that are missing.

The support packages are located in the `Product` directory of the Solaris CD-ROM. Use the `pkgadd` command to add any missing packages.

```
# pkgadd -d path_to_Solaris/Product Pkg1 Pkg2 Pkg3 ... PkgN
```

13. Shut down and power off Node A.

```
# shutdown -y -g0 -i0
```

For the procedure on shutting down and powering off a node, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

14. Install the host adapters in Node A.

For the procedure on installing host adapters, see the documentation that shipped with your host adapters and nodes.

15. Power on and boot Node A into non-cluster mode.

```
{0} ok boot -x
```

For more information, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

16. If necessary, upgrade the host adapter firmware on Node A.

See the *Sun Cluster 3.0 12/01 Release Notes* for information about accessing Sun's EarlyNotifier web pages, which list information about any required patches or firmware levels that are available for download. For the procedure on applying any host adapter firmware patch, see the firmware patch README file.

17. If necessary, install GBICs in the FC switches.

For the procedure on installing a GBIC to an FC switch, see the *SANbox 8/16 Segmented Loop Switch User's Manual*.

18. Connect fiber optic cables between Node A and the FC switches in your StorEdge 3900 Series system.

19. If necessary, install the required Solaris patches for StorEdge T3+ array support on Node A.

See the *Sun Cluster 3.0 12/01 Release Notes* for information about accessing Sun's EarlyNotifier web pages, which list information about any required patches or firmware levels that are available for download. For the procedure on applying any host adapter firmware patch, see the firmware patch README file.

20. Install any required patches or software for Sun StorEdge Traffic Manager software support to Node A from the Sun Download Center Web site,

<http://www.sun.com/storage/san/>

For instructions on installing the software, see the information on the web site.

21. Activate the Sun StorEdge Traffic Manager software functionality in the software you installed in [Step 20](#).

To activate the Sun StorEdge Traffic Manager software functionality, manually edit the `/kernel/drv/scsi_vhci.conf` file that is installed to change the `mpxio-disable` parameter to `no`:

```
mpxio-disable="no"
```

22. Shut down Node A.

```
# shutdown -y -g0 -i0
```

23. Perform a reconfiguration boot on Node A to create the new Solaris device files and links.

```
{0} ok boot -r
```


24. On Node A, update the /devices and /dev entries:

```
# devfsadm -C
```

25. On Node A, update the paths to the DID instances:

```
# scdidadm -C
```

26. Label the new array logical volume.

For the procedure on labeling a logical volume, see the??????????????????????.

27. (Optional) On Node A, verify that the device IDs (DIDs) are assigned to the new array.

```
# scdidadm -l
```

28. Do you need to install host adapters in Node B?

- If not, go to [Step 35](#).
- If you do need to install host adapters to Node B, continue with [Step 29](#).

29. Is the host adapter you are installing the first host adapter on Node B?

- If not, go to [Step 31](#).
- If it is the first host adapter, determine whether the required support packages for the host adapter are already installed on this node. The following packages are required.

```
# pkginfo | egrep Wlux
system    SUNWluxd   Sun Enterprise Network Array sf Device Driver
system    SUNWluxdx  Sun Enterprise Network Array sf Device Driver (64-bit)
system    SUNWluxl   Sun Enterprise Network Array social Device Driver
system    SUNWluxlx  Sun Enterprise Network Array social Device Driver (64-bit)
system    SUNWluxop  Sun Enterprise Network Array firmware and utilities
system    SUNWluxox  Sun Enterprise Network Array libraries (64-bit)
```

30. Are the required support packages already installed?

- If they are already installed, go to [Step 31](#).

- If not, install the missing support packages.

The support packages are located in the `Product` directory of the Solaris CD-ROM. Use the `pkgadd` command to add any missing packages.

```
# pkgadd -d path_to_Solaris/Product Pkg1 Pkg2 Pkg3 ... PkgN
```

31. Move all resource groups and device groups off Node B.

```
# scswitch -s -h nodename
```

32. Shut down and power off Node B.

```
# shutdown -y -g0 -i0
```

For the procedure on shutting down and powering off a node, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

33. Install the host adapters in Node B.

For the procedure on installing host adapters, see the documentation that shipped with your host adapters and nodes.

34. Power on and boot Node B.

```
{0} ok boot -x
```

For more information, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

35. If necessary, upgrade the host adapter firmware on Node B.

See the *Sun Cluster 3.0 12/01 Release Notes* for information about accessing Sun's EarlyNotifier web pages, which list information about any required patches or firmware levels that are available for download. For the procedure on applying any host adapter firmware patch, see the firmware patch `README` file.

36. If necessary, install GBICs in the FC switches.

For the procedure on installing a GBIC to an FC switch, see the *SANbox 8/16 Segmented Loop Switch User's Manual*.

37. Connect fiber optic cables between the FC switches in your StorEdge 3900 Series system and Node B.

38. **If necessary, install the required Solaris patches for StorEdge T3+ array support on Node B.**

For a list of required Solaris patches for StorEdge T3+ array support, see the *Sun StorEdge T3 Disk Tray Release Notes*.

39. **If you are installing a partner-group configuration, install any required patches or software for Sun StorEdge Traffic Manager software support to Node B from the Sun Download Center Web site, <http://www.sun.com/storage/san/>**

For instructions on installing the software, see the information on the web site.

40. **Activate the Sun StorEdge Traffic Manager software functionality in the software you installed in [Step 39](#).**

To activate the Sun StorEdge Traffic Manager software functionality, manually edit the `/kernel/drv/scsi_vhci.conf` file that is installed to change the `mpxio-disable` parameter to no:

```
mpxio-disable="no"
```

41. **Shut down Node B.**

```
# shutdown -y -g0 -i0
```

42. **Perform a reconfiguration boot to create the new Solaris device files and links on Node B.**

```
{0} ok boot -r
```

43. **On Node B, update the `/devices` and `/dev` entries:**

```
# devfsadm -C
```

44. **On Node B, update the paths to the DID instances:**

```
# scdidadm -C
```

45. **(Optional) On Node B, verify that the DIDs are assigned to the new arrays:**

```
# scdidadm -l
```

46. On one node attached to the new arrays, reset the SCSI reservation state:

```
# scdidadm -R n
```

Where *n* is the DID instance of a array LUN you are adding to the cluster.

Note – Repeat this command on the same node for each array LUN you are adding to the cluster.

47. Return the resource groups and device groups you identified in [Step 8](#) to all nodes.

```
# scswitch -z -g resource-group -h nodename  
# scswitch -z -D device-group-name -h nodename
```

For more information, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

48. Perform volume management administration to incorporate the new logical volumes into the cluster.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

▼ How to Remove StorEdge 3900 Series Systems From a Running Cluster

Use this procedure to permanently remove a StorEdge 3900 Series system and its associated submirrors from a running cluster.

This procedure defines “Node A” as the cluster node you begin working with, and “Node B” as the other node.



Caution – During this procedure, you lose access to the data that resides on each StorEdge T3+ array partner-group in the StorEdge 3900 Series system you are removing.

1. **If necessary, back up all database tables, data services, and volumes associated with each StorEdge T3+ partner-group in the StorEdge 3900 Series system you are removing.**
2. **If necessary, run the appropriate Solstice DiskSuite or VERITAS Volume Manager commands to detach the submirrors from each StorEdge T3+ partner-group in the StorEdge 3900 Series system that you are removing to stop all I/O activity to the partner-groups.**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

3. **Run the appropriate Solstice DiskSuite or VERITAS Volume Manager commands to remove references to each LUN that belongs to the StorEdge T3+ partner-group in the StorEdge 3900 Series system that you are removing.**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

4. **Determine the resource groups and device groups running on all nodes.**

Record this information because you will use it in [Step 16](#) of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

5. **Move all resource groups and device groups off Node A.**

```
# scswitch -s -h nodename
```

6. Shut down Node A.

```
# shutdown -y -g0 -i0
```

For the procedure on shutting down a node, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

7. Disconnect the cables that connected Node A to the FC switches in your StorEdge 3900 Series system.

8. Without allowing the node to boot, power on Node A.

For more information, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

9. Boot Node A into cluster mode.

```
{0} ok boot
```

10. Move all resource groups and device groups off Node B.

```
# scswitch -s -h nodename
```

11. Shut down Node B.

```
# shutdown -y -g0 -i0
```

12. Disconnect the cables that connected Node B to the FC switches in your StorEdge 3900 Series system.

13. Without allowing the node to boot, power on Node B.

For more information, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

14. Boot Node B into cluster mode.

```
{0} ok boot
```

For more information, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

15. On all cluster nodes, update the `/devices` and `/dev` entries.

```
# devfsadm -C
# sctidadm -C
```

16. Return the resource groups and device groups you identified in [Step 4](#) to all nodes.

```
# scswitch -z -g resource-group -h nodename
# scswitch -z -D device-group-name -h nodename
```

▼ How to Replace an FC Switch or Array-to-Switch Component in a Running Cluster

Use this procedure to replace an FC switch, or the following StorEdge T3+ array-to-switch components in a running cluster:

- Fiber optic cable that connects an FC switch to an array
- GBIC on an FC switch, connecting to an array
- StorEdge network FC switch-8 or switch-16
- Interconnect cables between two interconnected arrays of a partner-group

1. **Telnet to the StorEdge T3+ array that is connected to the FC switch or component that you are replacing.**
2. **Use the StorEdge T3+ `sys stat` command to view the controller status for the two arrays of the partner group.**

In the following example, both controllers are ONLINE.

```
t3://:<#> sys stat
Unit   State   Role   Partner
-----
 1     ONLINE   Master  2
 2     ONLINE   AlterM  1
```

See the *Sun StorEdge T3 and T3+ Array Administrator's Guide* for more information about the `sys stat` command.

3. **Is the FC switch or component that you are replacing attached to an array controller that is ONLINE or DISABLED, as determined in [Step 2](#)?**
 - If the controller is already DISABLED, go to [Step 5](#).
 - If the controller is ONLINE, use the T3+ `disable` command to disable it. Using the example from [Step 2](#), if you want to disable Unit 1, enter the following:

```
t3://:<#> disable u1
```

See the *Sun StorEdge T3 and T3+ Array Administrator's Guide* for more information about the `disable` command.

4. Use the T3+ `sys stat` command to verify that the controller's state has been changed to `DISABLED`.

```
t3://:<#> sys stat
Unit   State      Role      Partner
-----
1      DISABLED   Slave
2      ONLINE     Master
```

5. Replace the component using the following references:

- For the procedure on replacing a fiber optic cable between an array and an FC switch, see the *Sun StorEdge Network FC Switch-8 and Switch-16 Installation and Configuration Guide, Sun SAN 3.0*.
- For the procedure on replacing a GBIC on an FC switch, see the *SANbox 8/16 Segmented Loop Switch User's Manual*.
- For the procedure on replacing a StorEdge network FC switch-8 or switch-16, see the *Sun StorEdge 3900 and 6900 Series Reference Manual*.

Note – If you are replacing an FC switch and you intend to save the switch IP configuration for restoration to the replacement switch, do not connect the cables to the replacement switch until *after* you recall the Fabric configuration to the replacement switch. For more information about saving and recalling switch configurations see the *Sun StorEdge Network FC Switch-8 and Switch-16 Installation and Configuration Guide, Sun SAN 3.0*.

Note – Before you replace an FC switch, be sure that the `probe_timeout` parameter of your data service software is set to more than 90 seconds. Increasing the value of the `probe_timeout` parameter to more than 90 seconds avoids unnecessary resource group restarts when one of the FC switches is powered off.

- For the procedure on replacing interconnect cables, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.
6. If necessary, telnet to the array of the partner group that is still online.
 7. Use the StorEdge T3+ `enable` command to enable the array that you disabled in [Step 3](#).

```
t3://:<#> enable u1
```

See the *Sun StorEdge T3 and T3+ Array Administrator's Guide* for more information about the `enable` command.

8. Use the T3+ `sys stat` command to verify that the controller's state has been changed to ONLINE.

```
t3:/:<#> sys stat
Unit   State   Role   Partner
-----
 1     ONLINE   AlterM  2
 2     ONLINE   Master  1
```

▼ How to Replace a Node-to-Switch Component in a Running Cluster

Use this procedure to replace the following node-to-switch components in a running cluster:

- Node-to-switch fiber optic cable
- GBIC on an FC switch, connecting to a node

1. On the node connected to the component you are replacing, determine the resource groups and device groups running on the node.

Record this information because you will use it in [Step 4](#) of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

2. Move all resource groups and device groups to another node.

```
# scswitch -s -h nodename
```

3. Replace the node-to-switch component.

- For the procedure on replacing a fiber optic cable between a node and an FC switch, see the *Sun StorEdge network FC switch-8 and switch-16 Installation and Configuration Guide*.
- For the procedure on replacing a GBIC on an FC switch, see the *SANbox 8/16 Segmented Loop Switch User's Manual*.

4. Return the resource groups and device groups you identified in [Step 1](#) to the node that is connected to the component you replaced.

```
# scswitch -z -g resource-group -h nodename  
# scswitch -z -D device-group-name -h nodename
```

▼ How to Remove StorEdge T3/T3+ Arrays From a Running Cluster

Use this procedure to permanently remove StorEdge T3/T3+ array partner groups and their submirrors from a running cluster.

This procedure defines “Node A” as the cluster node you begin working with, and “Node B” as the other node.



Caution – During this procedure, you lose access to the data that resides on each StorEdge T3+ array partner-group you are removing.

1. **If necessary, back up all database tables, data services, and volumes associated with each partner-group you are removing.**
2. **If necessary, run the appropriate Solstice DiskSuite or VERITAS Volume Manager commands to detach the submirrors from each array or partner-group that you are removing to stop all I/O activity to the array or partner-group.**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

3. **Run the appropriate Solstice DiskSuite or VERITAS Volume Manager commands to remove references to each LUN that belongs to the array or partner-group that you are removing.**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

4. **Determine the resource groups and device groups running on all nodes.**

Record this information because you will use it in [Step 21](#) of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

5. **Move all resource groups and device groups off Node A.**

```
# scswitch -S -h nodename
```

6. Shut down Node A.

```
# shutdown -y -g0 -i0
```

For the procedure on shutting down a node, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

7. Disconnect from both arrays the fiber optic cables connecting to the FC switches, then the Ethernet cable(s).

8. Is any array you are removing the last array connected to an FC switch on Node A?

- If not, go to [Step 12](#).
- If it is the last array, disconnect the fiber optic cable between Node A and the FC switch that was connected to this array.

For the procedure on removing a fiber optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

9. Do you want to remove the host adapters from Node A?

- If not, go to [Step 12](#).
- If yes, power off Node A.

10. Remove the host adapters from Node A.

For the procedure on removing host adapters, see the documentation that shipped with your host adapter and nodes.

11. Without allowing the node to boot, power on Node A.

For more information, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

12. Boot Node A into cluster mode.

```
{0} ok boot
```

13. Move all resource groups and device groups off Node B.

```
# scswitch -s -h nodename
```

14. Shut down Node B.

```
# shutdown -y -g0 -i0
```

15. Is any array you are removing the last array connected to an FC switch on Node B?

- If not, go to [Step 19](#).
- If it is the last array, disconnect the fiber optic cable connecting this FC switch to Node B.

For the procedure on removing a fiber optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

16. Do you want to remove the host adapters from Node B?

- If not, go to [Step 19](#).
- If yes, power off Node B.

17. Remove the host adapters from Node B.

For the procedure on removing host adapters, see the documentation that shipped with your nodes.

18. Without allowing the node to boot, power on Node B.

For more information, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

19. Boot Node B into cluster mode.

```
{0} ok boot
```

For more information, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

20. On all cluster nodes, update the `/devices` and `/dev` entries.

```
# devfsadm -C
# sctdidadm -C
```

21. Return the resource groups and device groups you identified in [Step 4](#) to all nodes.

```
# scswitch -z -g resource-group -h nodename
# scswitch -z -D device-group-name -h nodename
```

▼ How to Upgrade StorEdge T3+ Array Firmware

Use one of the following procedures to upgrade the firmware on the StorEdge T3+ arrays in your StorEdge 3900 Series system, depending on whether your array partner-group has been configured to support submirrors of a cluster node's volumes. StorEdge T3+ array firmware includes controller firmware, unit interconnect card (UIC) firmware, EPROM firmware, and disk drive firmware.

- [“Upgrading Firmware on Arrays That Support Submirrored Data” on page 95](#)
- [“Upgrading Firmware on Arrays That Do Not Support Submirrored Data” on page 97](#)

Note – For all firmware, always read any README files that accompany the firmware for the latest information and special notes.

Upgrading Firmware on Arrays That Support Submirrored Data



Caution – Perform this procedure on one array at a time. This procedure requires that you reset the arrays you are upgrading. If you reset more than one array at a time, your cluster will lose access to data.

- 1. On the node that currently owns the disk group or disk set to which the submirror belongs, detach the submirrors of the array on which you are upgrading firmware. (This procedure refers to this node as Node A and remaining node as Node B.)**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

- 2. Disconnect both array-to-switch fiber optic cables from the two arrays of the partner-group.**
- 3. Apply the controller, disk drive, and UIC firmware patches.**

For the list of required StorEdge T3+ array patches, see the *Sun StorEdge T3 Disk Tray Release Notes*. For the procedure on applying firmware patches, see the firmware patch README file. For the procedure on verifying the firmware level, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

- 4. Reset the arrays.**

For the procedure on resetting an array, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

5. Use the StorEdge T3+ `disable` command to disable the array controller that is attached to Node B so that all logical volumes come under the control of the remaining controller.

```
t3:/:<#> disable uencidctr
```

See the *Sun StorEdge T3 and T3+ Array Administrator's Guide* for more information about the `disable` command.

6. Reconnect both array-to-switch fiber optic cables to the two arrays of the partner-group.
7. On one node connected to the partner-group, use the `format` command to verify that the array controllers are rediscovered by the node.

```
# format
```

8. Use the StorEdge T3+ `enable` command to enable the array controller that you disabled in [Step 5](#).

```
t3:/:<#> enable uencidctr
```

9. Reattach the submirrors that you detached in [Step 1](#) to resynchronize them.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

Upgrading Firmware on Arrays That Do Not Support Submirrored Data

In a partner-pair configuration, it is possible to have non-mirrored data; however, this requires that you shutdown the cluster when upgrading firmware, as described in this procedure.

1. Shutdown the entire cluster.

```
# scshutdown -y -g0
```

For the full procedure on shutting down a cluster, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

2. Apply the controller, disk drive, and UIC firmware patches.

For the list of required StorEdge T3+ array patches, see the *Sun StorEdge T3 Disk Tray Release Notes*. For the procedure on applying firmware patches, see the firmware patch README file. For the procedure on verifying the firmware level, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

3. Reset the arrays.

For the procedure on resetting an array, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

4. Boot all nodes back into the cluster.

```
ok boot
```

For the full procedure on booting nodes into the cluster, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

5. On one node connected to the partner-group, use the `format` command to verify that the array controllers are rediscovered by the node.

```
# format
```

▼ How to Replace a Failed Disk Drive in a Running Cluster

Use this procedure to replace one failed disk drive in a StorEdge T3+ array that is in your StorEdge 3900 Series system, in a running cluster.



Caution – If you remove any field replaceable unit (FRU) for an extended period of time, thermal complications might result. To prevent this complication, the StorEdge T3+ array is designed so that an orderly shutdown occurs when you remove a component for longer than 30 minutes. Therefore, a replacement part must be immediately available before starting an FRU replacement procedure. You must replace an FRU within 30 minutes or the StorEdge T3+ array, and all attached StorEdge T3+ arrays, will shut down and power off.

1. Did the failed disk drive impact the array logical volume's availability?

- If not, go to [Step 2](#).
- If it did impact logical volume availability, remove the logical volume from volume management control.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

2. Replace the disk drive in the array.

For the procedure on replacing a disk drive, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

3. Did you remove a LUN from volume management control in [Step 1](#)?

- If not, you are finished with this procedure.
- If you did remove a LUN from volume management control, return the LUN to volume management control now.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

▼ How to Replace a Node's Host Adapter in a Running Cluster

Use this procedure to replace a failed host adapter in a running cluster. As defined in this procedure, "Node A" is the node with the failed host adapter you are replacing and "Node B" is the other node.

1. Determine the resource groups and device groups running on all nodes.

Record this information because you will use it in [Step 8](#) of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

2. Move all resource groups and device groups off Node A.

```
# scswitch -s -h nodename
```

3. Shut down Node A.

```
# shutdown -y -g0 -i0
```

4. Power off Node A.

For more information, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

5. Replace the failed host adapter.

For the procedure on removing and adding host adapters, see the documentation that shipped with your nodes.

6. Power on Node A.

For more information, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

7. Boot Node A into cluster mode.

```
{0} ok boot
```

For more information, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

8. Return the resource groups and device groups you identified in [Step 1](#) to all nodes.

```
# scswitch -z -g resource-group -h nodename  
# scswitch -z -D device-group-name -h nodename
```

For more information, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

Installing and Maintaining a Sun StorEdge 9910 or StorEdge 9960 Array

This chapter contains the procedures for installing, configuring, and maintaining Sun StorEdge 9910 and Sun StorEdge 9960 arrays.

This chapter contains the following procedures:

- [“How to Add a StorEdge 9910/9960 Array Logical Volume to a Cluster” on page 102](#)
- [“How to Remove a StorEdge 9910/9960 Array Logical Volume” on page 103](#)
- [“How to Add a StorEdge 9910/9960 Array” on page 106](#)
- [“How to Remove a StorEdge 9910/9960 Array” on page 110](#)
- [“How to Replace a Host-to-Array Fiber-Optic Cable” on page 113](#)
- [“How to Replace a Host Adapter” on page 113](#)

For conceptual information on multihost disks, see the *Sun Cluster 3.0 12/01 Concepts* document.

Configuring a StorEdge 9910/9960 Array

This section contains the procedures for configuring a StorEdge 9910/9960 array in a running cluster. The following table lists these procedures.

TABLE D-1 Task Map: Configuring a StorEdge 9910/9960 Array

Task	For Instructions, Go To
Add a StorEdge 9910/9960 array logical volume	See “How to Add a StorEdge 9910/9960 Array Logical Volume to a Cluster” on page 102.
Remove a StorEdge 9910/9960 array logical volume	See “How to Remove a StorEdge 9910/9960 Array Logical Volume” on page 103.

▼ How to Add a StorEdge 9910/9960 Array Logical Volume to a Cluster

Use this procedure to add a logical volume to a cluster. This procedure assumes that your service provider has created your logical volumes and that all cluster nodes are booted and attached to the StorEdge 9910/9960 array.

1. **On all cluster nodes, update the `/devices` and `/dev` entries.**

```
# devfsadm
```

2. **On one node connected to the partner-group, use the `format` command to verify that the new logical volume is visible to the system.**

```
# format
```

See the `format` command man page for more information about using the command.

3. Are you running VERITAS Volume Manager?

- If not, go to [Step 4](#)
- If you are running VERITAS Volume Manager, update its list of devices on all cluster nodes attached to the logical volume that you created in [Step 2](#).

See your VERITAS Volume Manager documentation for information about using the `vxdtl enable` command to update new devices (volumes) in your VERITAS Volume Manager list of devices.

4. If necessary, partition the logical volume.

Contact your service provider to partition the logical volume.

5. From any node in the cluster, update the global device namespace.

```
# scgdevs
```

If a volume management daemon such as `vold` is running on your node, and you have a CD-ROM drive that is connected to the node, a `device busy` error might be returned even if no disk is in the drive. This error is expected behavior.

Where to Go From Here

To create a new resource or reconfigure a running resource to use the new StorEdge 9910/9960 Array logical volume, see the *Sun Cluster 3.0 12/01 Data Services Installation and Configuration Guide*.

▼ How to Remove a StorEdge 9910/9960 Array Logical Volume

Use this procedure to remove a logical volume. This procedure assumes all cluster nodes are booted and attached to the StorEdge 9910/9960 array that hosts the logical volume you are removing.

This procedure defines Node A as the node you begin working with, and Node B as the remaining node.



Caution – This procedure removes all data on the logical volume you are removing.

1. If necessary, migrate all data and volumes off the logical volume that you are removing. Otherwise, proceed to [Step 2](#).

2. Are you running VERITAS Volume Manager?

- If no, go to [Step 3](#).
- If yes, update the list of devices on all cluster nodes attached to the logical volume that you are removing.

See your VERITAS Volume Manager documentation for information about using the `vxdisk rm` command to remove devices (volumes) in your VERITAS Volume Manager device list.

3. Run the appropriate Solstice DiskSuite or VERITAS Volume Manager commands to remove the reference to the logical unit number (LUN) from any diskset or disk group.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

4. Remove the logical volume.

Contact your service provider to remove the logical volume.

5. Determine the resource groups and device groups that are running on Node A and Node B.

Record this information because you will use it in [Step 12](#) of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

6. Move all resource groups and device groups off Node A.

```
# scswitch -s -h from-node
```

7. Shut down and reboot Node A by using the `shutdown` command with the `-i6` option.

The `-i6` option with the `shutdown` command causes the node to shutdown and reboot.

```
# shutdown -y -g0 -i6
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

8. On Node A, update the `/devices` and `/dev` entries.

```
# devfsadm -C
# sctidadm -C
```

9. Move all resource groups and device groups off Node B.

```
# scswitch -s -h from-node
```

10. Shut down and reboot Node B by using the `shutdown` command with the `-i6` option.

The `-i6` option with the `shutdown` command causes the node to reboot after it shuts down to the `ok` prompt.

```
# shutdown -y -g0 -i6
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

11. On Node B, update the `/devices` and `/dev` entries.

```
# devfsadm -C
# sctidadm -C
```

12. Return the resource groups and device groups you identified in [Step 5](#) to Node B.

```
# scswitch -z -g resource-group -h nodename
# scswitch -z -D device-group-name -h nodename
```

For more information, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

Where to Go From Here

To create a logical volume, see [“How to Add a StorEdge 9910/9960 Array Logical Volume to a Cluster”](#) on page 102.

Maintaining a StorEdge 9910/9960 Array

This section contains a limited set of procedures for maintaining a StorEdge 9910/9960 array. Contact your service provider to add, remove, or replace any StorEdge 9910/9960 components.

TABLE D-2 Task Map: Maintaining a StorEdge 9910/9960 Array

Task	For Instructions, Go To
Add a StorEdge 9910/9960 array.	"How to Add a StorEdge 9910/9960 Array" on page 106
Remove a StorEdge 9910/9960 array.	"How to Remove a StorEdge 9910/9960 Array" on page 110
Replace a host-to-array fiber-optic cable.	"How to Replace a Host-to-Array Fiber-Optic Cable" on page 113
Replace a host adapter.	"How to Replace a Host Adapter" on page 113

▼ How to Add a StorEdge 9910/9960 Array

Use this procedure to add a new StorEdge 9910/9960 array to a running cluster.

This procedure defines Node A as the node you begin working with, and Node B as the remaining node.

1. Power on the StorEdge 9910/9960 array.

Note – The StorEdge 9910/9960 array will require a few minutes to boot.

Contact your service provider to power on the StorEdge 9910/9960 array.

2. Configure the new StorEdge 9910/9960 array.

Contact your service provider to create the desired logical volumes.

3. Determine the resource groups and device groups that are running on Node A and Node B.

Record this information because you will use it in [Step 28](#) of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

4. Move all resource groups and device groups off Node A.

```
# scswitch -s -h from-node
```

5. Do you need to install a host adapter in Node A?

- If no, skip to [Step 11](#).
- If yes, proceed to [Step 6](#).

6. Is the host adapter the first JNI host adapter on Node A?

- If no, skip to [Step 11](#).
- If yes, contact your service provider to install the support packages and configure the drivers.

7. Stop the Sun Cluster software on Node A and shut down Node A.

```
# shutdown -y -g0 -i0
```

For the procedure on shutting down a node, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

8. Power off Node A.

9. Install the host adapter in Node A.

For the procedure on installing a host adapter, see the documentation that shipped with your host adapter or updated information on the manufacturer's web site.

10. If necessary, power on and boot Node A.

```
{0} ok boot -x
```

For more information, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

11. Attach the StorEdge 9910/9960 array to Node A.

Contact your service provider to install a fiber-optic cable between the StorEdge 9910/9960 array and your cluster node.

12. Configure the fibre-channel adapter and the StorEdge 9910/9960 array.

Contact your service provider to configure the adapter and StorEdge 9910/9960 array.

13. Shut down Node A.

```
# shutdown -y -g0 -i0
```

14. Perform a reconfiguration boot to create the new Solaris device files and links on Node A.

```
{0} ok boot -r
```

15. Label the new logical volume.

Contact your service provider to label the new logical volume.

16. (Optional) On Node A, verify that the device IDs (DIDs) are assigned to the new StorEdge 9910/9960 array.

```
# scdidadm -l
```

17. Do you need to install a host adapter in Node B?

- If yes, proceed to [Step 18](#).
- If no, skip to [Step 24](#).

18. Is the host adapter the first JNI host adapter on Node B?

- If no, skip to [Step 19](#).
- If yes, contact your service provider to install the support packages and configure the drivers.

19. Move all resource groups and device groups off Node B.

```
# scswitch -s -h from-node
```

20. Stop the Sun Cluster software on Node B, and shut down the node.

```
# shutdown -y -g0 -i0
```

For the procedure on shutting down a node, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

21. Power off Node B.

For more information, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

22. Install the host adapter in Node B.

For the procedure on installing a host adapter, see the documentation that shipped with your host adapter and node.

23. If necessary, power on and boot Node B.

```
{0} ok boot -x
```

For more information, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

24. Attach the StorEdge 9910/9960 array to Node B.

Contact your service provider to install a fiber-optic cable between the StorEdge 9910/9960 array and your cluster node.

25. Shut down Node B.

```
# shutdown -y -g0 -i0
```

26. Perform a reconfiguration boot to create the new Solaris device files and links on Node B.

```
{0} ok boot -r
```

27. (Optional) On Node B, verify that the device IDs (DIDs) are assigned to the new StorEdge 9910/9960 array.

```
# scdidadm -l
```

28. Return the resource groups and device groups you identified in [Step 3](#) to Node B.

```
# scswitch -z -g resource-group -h nodename
# scswitch -z -D device-group-name -h nodename
```

For more information, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

29. **Perform volume management administration to incorporate the new logical volumes into the cluster.**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

▼ How to Remove a StorEdge 9910/9960 Array

Use this procedure to permanently remove a StorEdge 9910/9960 array and its submirrors from a running cluster. This procedure provides the flexibility to remove the host adapters from the nodes for the StorEdge 9910/9960 array you are removing.

This procedure defines Node A as the node you begin working with, and Node B as the remaining node.



Caution – During this procedure, you will lose access to the data that resides on the StorEdge 9910/9960 array you are removing.

1. **Back up all database tables, data services, and volumes that are associated with the StorEdge 9910/9960 array that you are removing.**
2. **Detach the submirrors from the StorEdge 9910/9960 array you are removing in order to stop all I/O activity to the array.**
For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.
3. **Run the appropriate Solstice DiskSuite or VERITAS Volume Manager commands to remove the references to the LUN(s) from any diskset or disk group.**
For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.
4. **Determine the resource groups and device groups that are running on Node A.**

```
# scstat
```

5. Move all resource groups and device groups off Node A.

```
# scswitch -s -h from-node
```

6. Stop the Sun Cluster software on Node A, and shut down Node A.

```
# shutdown -y -g0 -i0
```

For the procedure on shutting down a node, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

7. Disconnect the fiber-optic cable between Node A and the StorEdge 9910/9960 array.

8. Do you want to remove the host adapter from Node A?

- If no, skip to [Step 11](#).
- If yes, power off Node A.

9. Remove the host adapter from Node A.

For the procedure on removing host adapters, see the documentation that shipped with your host adapter or updated information on the manufacturer's web site.

10. Without allowing the node to boot, power on Node A.

For more information, see the documentation that shipped with your server.

11. Boot Node A into cluster mode.

```
{0} ok boot
```

12. Determine the resource groups and device groups that are running on Node A.

```
# scstat
```

13. Move all resource groups and device groups off Node B.

```
# scswitch -s -h from-node
```

14. Stop the Sun Cluster software on Node B, and shut down Node B.

```
# shutdown -y -g0 -i0
```

15. Disconnect the fiber-optic cable between Node B and the StorEdge 9910/9960 array.
16. Do you want to remove the host adapter from Node B?
 - If no, skip to [Step 19](#).
 - If yes, power off Node B.

17. Remove the host adapter from Node B.

For the procedure on removing host adapters, see the documentation that shipped with your server and host adapter.

18. Without allowing the node to boot, power on Node B.

For more information, see the documentation that shipped with your server.

19. Boot Node B into cluster mode.

```
{0} ok boot
```

For more information, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

20. On all cluster nodes, update the `/devices` and `/dev` entries.

```
# devfsadm -C  
# sctdidadm -C
```

21. Return the resource groups and device groups you identified in [Step 4](#) to Node B.

```
# scswitch -z -g resource-group -h nodename  
# scswitch -z -D device-group-name -h nodename
```


▼ How to Replace a Host-to-Array Fiber-Optic Cable

Use this procedure to replace the host-to-array fiber-optic cables. Node A in this procedure refers to the node with the failed fiber-optic cable that you are replacing.

1. **On Node A, determine the resource groups and device groups that are running on this node.**

Record this information because you will use it in [Step 4](#) of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

2. **Move all resource groups and device groups off Node A.**

```
# scswitch -s -h from-node
```

3. **Replace the host-to-array fiber-optic cable.**
4. **Return the resource groups and device groups you identified in [Step 1](#) to Node A.**

```
# scswitch -z -g resource-group -h nodename  
# scswitch -z -D device-group-name -h nodename
```

▼ How to Replace a Host Adapter

Use this procedure to replace a failed host adapter in a running cluster. Node A in this procedure refers to the node with the failed host adapter you are replacing. Node B is a backup node.

1. **Determine the resource groups and device groups that are running on Node A and Node B.**

Record this information because you will use it in [Step 8](#) of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

2. Move all resource groups and device groups off Node A.

```
# scswitch -s -h from-node
```

3. Shut down Node A.

```
# shutdown -y -g0 -i0
```

4. Power off Node A.

For more information, see the documentation that shipped with your server.

5. Replace the failed host adapter.

Contact your service provider to replace the failed host adapter.

6. Power on Node A.

For more information, see the documentation that shipped with your server.

7. Boot Node A into cluster mode.

```
{0} ok boot
```

For more information, see the *Sun Cluster 3.0 12/01 System Administration Guide*.

8. Return the resource groups and device groups you identified in [Step 1](#) to Node A.

```
# scswitch -z -g resource-group -h nodename  
# scswitch -z -D device-group-name -h nodename
```

For more information, see the *Sun Cluster 3.0 12/01 System Administration Guide*.