



Sun Cluster 3.x Release Notes Supplement

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 816-3381-10
July 2003, Revision A

Copyright 2003 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

This product or document is protected by copyright and 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, docs.sun.com, AnswerBook, AnswerBook2, Solstice DiskSuite, and Solaris are trademarks, registered trademarks, or service marks 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 an architecture developed by Sun Microsystems, Inc. This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>).

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

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

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 2003 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Ce produit ou document est protégé par un copyright et 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 du système 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, docs.sun.com, AnswerBook, AnswerBook2, Solstice DiskSuite, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, 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. Ce produit inclut le logiciel développé par la base de Apache Software Foundation (<http://www.apache.org/>).

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



030711@5943



Contents

1 Sun Cluster 3.1 Release Notes Supplement	9
Revision Record	9
New Features	10
Restrictions and Requirements	10
Changing Quorum Device Connectivity	11
Required VxFS Default Stack Size Increase	11
Known Problems	11
Bug ID 4840853	11
Bug ID 4863254	11
Bug ID 4867560	12
Bug ID 4867584	12
Known Documentation Problems	12
2 Sun Cluster Data Services 3.1 5/03 Release Notes Supplement	13
Revision Record	13
New Features	14
Sun Cluster Support for Oracle Parallel Server/Real Application Clusters on a Subset of Cluster Nodes	14
Restrictions and Requirements	15
Known Problems	15
Known Documentation Problems	16
Sun Cluster 3.1 Data Service for Sun ONE Application Server	16
3 Sun Cluster 3.0 5/02 Release Notes Supplement	17
Revision Record	17

New Features	25
Security Hardening for Solaris 9	25
Failover File System (HAStoragePlus)	25
RAID 5 on Sun StorEdge 99x0 Storage Arrays	26
Apache 2.0	26
New Guidelines for the swap Partition	26
Support for Oracle Parallel Server/Real Application Clusters on the Cluster File System	26
Restrictions and Requirements	28
Dynamic Multipathing (DMP)	28
Changing Quorum Device Connectivity	29
Storage Topologies Replaced by New Requirements	29
Shared Storage Restriction Relaxed	29
EMC Storage Restriction	29
Framework Restrictions and Requirements	30
Oracle UDLM Requirement	30
Known Problems	31
Bug ID 4346123	31
Bug ID 4662264	31
Bug ID 4665886	32
Bug ID 4668496	32
Bug ID 4680862	32
Bug ID 4779686	32
BugId 4818874	33
BugId 4836405	33
BugID 4838619	34
Known Documentation Problems	34
Hardware Guide	34
Software Installation Guide	36
Data Services Installation and Configuration Guide	43
Supplement	52
Release Notes	53
Man Pages	53
A Scalable Cluster Topology	57
Overview of Scalable Topology	57
Adding or Removing a Cluster Node	58
Adding a Cluster Node	58

	Removing a Cluster Node	58
	▼ How to Remove Connectivity Between an Array and a Single Node, in a Cluster With Greater Than Two-Node Connectivity	60
B	Installing and Configuring Sun Cluster HA for SAP liveCache	63
	Sun Cluster HA for SAP liveCache Overview	63
	Installing and Configuring Sun Cluster HA for SAP liveCache	65
	Planning the Sun Cluster HA for SAP liveCache Installation and Configuration	66
	Configuration Requirements	67
	Standard Data Service Configurations	67
	Configuration Considerations	68
	Configuration Planning Questions	68
	Preparing the Nodes and Disks	69
	▼ How to Prepare the Nodes	69
	Installing and Configuring liveCache	69
	▼ How to Install and Configure liveCache	70
	▼ How to Enable liveCache to Run in a Cluster	70
	Verifying the liveCache Installation and Configuration	71
	▼ How to Verify the liveCache Installation and Configuration	71
	Installing the Sun Cluster HA for SAP liveCache Packages	72
	▼ How to Install the Sun Cluster HA for SAP liveCache Packages	72
	Registering and Configuring the Sun Cluster HA for SAP liveCache	73
	Sun Cluster HA for SAP liveCache Extension Properties	73
	▼ How to Register and Configure Sun Cluster HA for SAP liveCache	75
	Verifying the Sun Cluster HA for SAP liveCache Installation and Configuration	78
	▼ How to Verify the Sun Cluster HA for SAP liveCache Installation and Configuration	78
	Understanding Sun Cluster HA for SAP liveCache Fault Monitors	80
	Extension Properties	80
	Monitor Check Method	80
	Probing Algorithm and Functionality	80
C	Installing and Configuring Sun Cluster HA for Sybase ASE	83
	Installing and Configuring Sun Cluster HA for Sybase ASE	83
	Preparing to Install Sun Cluster HA for Sybase ASE	84
	Installing the Sybase ASE 12.0 Software	85
	▼ How to Prepare the Nodes	85

- ▼ How to Install the Sybase Software 87
- ▼ How to Verify the Sybase ASE Installation 88
- Creating the Sybase ASE Database Environment 89
 - ▼ How to Configure Sybase ASE Database Access With Solstice DiskSuite/Solaris Volume Manager 89
 - ▼ How to Configure Sybase ASE Database Access With VERITAS Volume Manager 90
 - ▼ How to Create the Sybase ASE Database Environment 91
- Installing the Sun Cluster HA for Sybase ASE Package 92
 - ▼ How to Install Sun Cluster HA for Sybase ASE Packages 93
- Registering and Configuring Sun Cluster HA for Sybase ASE 93
 - ▼ How to Register and Configure Sun Cluster HA for Sybase ASE 94
- Verifying the Sun Cluster HA for Sybase ASE Installation 96
 - ▼ How to Verify the Sun Cluster HA for Sybase ASE Installation 97
- Understanding Sun Cluster HA for Sybase ASE Logging and Security Issues 98
 - Sun Cluster HA for Sybase ASE Logging 98
 - Important Security Issues 98
- Configuring Sun Cluster HA for Sybase ASE Extension Properties 99
- Sun Cluster HA for Sybase ASE Fault Monitor 101
 - Main Fault Monitor Process 102
 - Database-Client Fault Probe 102
 - Extension Properties 103

D RSM Phase II: RSMRDT Driver Installation 105

- Overview of the RSMRDT Driver 105
 - Installing the RSMRDT Driver 106
 - Restrictions 106
 - ▼ How to Install the SUNWscrdt Package 106
 - ▼ How to Uninstall the SUNWscrdt Package 106
 - ▼ How to Unload the RSMRDT Driver Manually 107

E Installing and Configuring Sun Cluster HA for SAP 109

- Sun Cluster HA for SAP Overview 110
- Installing and Configuring Sun Cluster HA for SAP 111
- Planning the Sun Cluster HA for SAP Installation and Configuration 112
 - Configuration Restrictions 112
 - Configuration Requirements 113

Standard Data Service Configurations	114
Configuration Considerations	115
Configuration Planning Questions	116
Packages and Support	117
Upgrading Sun Cluster HA for SAP	117
▼ How to Upgrade a Resource Type or Convert a Failover Application Resource to a Scalable Application Resource	118
Preparing the Nodes and Disks	119
▼ How to Prepare the Nodes	119
Installing and Configuring SAP and Database	119
▼ How to Install SAP and the Database	120
▼ How to Install an SAP Scalable Application Server	120
▼ How to Enable Failover SAP Instances to Run in a Cluster	123
Configuring Sun Cluster HA for DBMS	125
Where to Go From Here	125
Verifying the SAP Installation	125
▼ How to Verify SAP and the Database Installation with Central Instance	125
▼ How to Verify an SAP Failover Application Server	126
How to Verify an SAP Scalable Application Server	128
Installing the Sun Cluster HA for SAP Packages	128
▼ How to Install the Sun Cluster HA for SAP Packages	128
Registering and Configuring Sun Cluster HA for SAP	129
Sun Cluster HA for SAP Extension Properties	129
▼ How to Register and Configure Sun Cluster HA for SAP with Central Instance	136
▼ How to Register and Configure Sun Cluster HA for SAP as a Failover Data Service	137
▼ How to Register and Configure Sun Cluster HA for SAP as a Scalable Data Service	138
Setting Up a Lock File	139
▼ How to Set Up a Lock File for Central Instance or the Failover Application Server	140
▼ How to Set Up a Lock File for Scalable Application Server	141
Verifying the Sun Cluster HA for SAP Installation and Configuration	142
▼ How to Verify Sun Cluster HA for SAP Installation and Configuration and Central Instance	142
▼ How to Verify the Installation and Configuration of Sun Cluster HA for SAP as a Failover Data Service	143
▼ How to Verify Sun Cluster HA for SAP Installation and Configuration of as a Scalable Data Service	143

Understanding Sun Cluster HA for SAP Fault Monitor	144
Sun Cluster HA for SAP Fault Probes for Central Instance	144
Sun Cluster HA for SAP Fault Probes for Application Server	145
F Upgrading Sun Cluster Software	147
Upgrading to Sun Cluster 3.1 Software	147
Upgrade Requirements and Restrictions	148
▼ How to Prepare the Cluster for Upgrade	149
▼ How to Upgrade the Solaris Operating Environment	152
▼ How to Upgrade to Sun Cluster 3.1 Software	154
▼ How to Upgrade Sun Cluster-Module Software for Sun Management Center	160
▼ How to Finish Upgrading to Sun Cluster 3.1 Software	161
Recovering From Storage Configuration Changes During Upgrade	162
▼ How to Handle Storage Reconfiguration During an Upgrade	162
▼ How to Resolve Mistaken Storage Changes During an Upgrade	163

Sun Cluster 3.1 Release Notes Supplement

This chapter supplements the standard user documentation, including the *Sun Cluster 3.1 Release Notes* shipped with the Sun™ Cluster 3.1 product. These “online release notes” provide the most current information on the Sun Cluster 3.1 product. This chapter includes the following information.

- “Revision Record” on page 9
- “New Features” on page 10
- “Restrictions and Requirements” on page 10
- “Known Problems” on page 11
- “Known Documentation Problems” on page 12

Revision Record

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

TABLE 1-1 Sun Cluster 3.1 Release Notes Supplement Revision Record: Year 2003

Revision Date	New Information
July 2003	Procedures to enable Sun Cluster Support for Oracle Parallel Server/Real Application Clusters on a subset of cluster nodes. See “Sun Cluster Support for Oracle Parallel Server/Real Application Clusters on a Subset of Cluster Nodes” on page 14. Revised support for Multiple Masters configuration of Sun Cluster HA for Sun ONE Application Server. See “Sun Cluster 3.1 Data Service for Sun ONE Application Server” on page 16.

TABLE 1–1 Sun Cluster 3.1 Release Notes Supplement Revision Record: Year 2003
(Continued)

Revision Date	New Information
June 2003	<p>Procedures to upgrade a Sun Cluster 3.0 configuration to Sun Cluster 3.1 software, including upgrading from Solaris 8 to Solaris 9 software. See Appendix F.</p> <p>Modifications to make to the <code>/etc/system</code> file to correct changes made by VxFS installation. See “Changing Quorum Device Connectivity” on page 11.</p> <p>Procedures to support Sun StorEdge 6320 storage systems. See “Installing and Maintaining a Sun StorEdge 6320 System” in <i>Sun Cluster 3.x With Sun StorEdge 6320 System Manual</i>.</p> <p>Sun StorEdge 6120 storage arrays in dual-controller configurations and Sun StorEdge 6320 storage systems are limited to four nodes and 16 LUNs. See “Bug ID 4840853” on page 11.</p> <p>Procedures to support Sun StorEdge 3510 FC storage device. See the <i>Sun Cluster 3.x With Sun StorEdge 3510 FC Array Manual</i>.</p> <p>Sun StorEdge 3510 FC storage arrays are limited to 256 LUNs per channel. See “Bug ID 4867584” on page 12.</p> <p>Sun StorEdge 3510 FC storage arrays are limited to one node per channel. See “Bug ID 4867560” on page 12.</p>

New Features

In addition to features documented in *Sun Cluster 3.1 Release Notes*, this release now includes support for the following features.

There are no new features at this time.

Restrictions and Requirements

The following restrictions and requirements have been added or updated since the Sun Cluster 3.1 release.

Changing Quorum Device Connectivity

When you increase or decrease the number of node attachments to a quorum device, the quorum vote count is not automatically recalculated. You can reestablish the correct quorum vote if you remove all quorum devices and then add them back into the configuration.

Required VxFS Default Stack Size Increase

The default stack size that VERITAS File System (VxFS) sets during installation is 0x4000. However, this is inadequate for Sun Cluster software and might lead to a system panic. If you install VxFS on a Sun Cluster configuration, you must reset the stack size by making the following modifications to entries in the `/etc/system` file on each cluster node.

```
set rpcmod:svc_default_stksize=0x8000
set lwp_default_stksize=0x6000
```

- The first line changes the value for the `rpcmod:svc_default_stksize` variable from 0x4000 to 0x8000.
- The second line sets the value of the `lwp_default_stksize` variable to 0x6000.

Known Problems

In addition to known problems documented in *Sun Cluster 3.1 Release Notes*, the following known problems affect the operation of the Sun Cluster 3.1 release.

Bug ID 4840853

Problem Summary: Due to memory segmentation issues, if you configure the StorEdge 6120 or StorEdge 6320 storage system with four nodes and more than 16 LUNs, the storage device might fail and cause your data to be compromised.

Workaround: When using a StorEdge 6120 or StorEdge 6320 storage system, you must limit your configuration to four nodes and 16 LUNs.

Bug ID 4863254

Problem Summary: Due to a Solaris bug (4511634), Sun Cluster 3.1 does not provide the ability to auto-create IPMP groups when you add a logical host.

Workaround: You must manually create an IPMP group when you add a logical host.

Bug ID 4867560

Problem Summary: When two nodes are connected to the same channel of a Sun StorEdge 3510 FC storage array, rebooting one node causes the other node to lose the SCSI-2 reservation.

Workaround: You can only connect one node per channel on the Sun StorEdge 3510 FC storage arrays.

Bug ID 4867584

Problem Summary: If you have 512 LUNs in a direct-attach storage configuration with Sun StorEdge 3510 FC storage arrays, LUNs might be lost when the server reboots.

Workaround: Sun StorEdge 3510 FC storage arrays are limited to 256 LUNs per channel (eight IDs) in a Sun Cluster configuration.

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 3.1 Release Notes*.

There are no known problems at this time.

Sun Cluster Data Services 3.1 5/03 Release Notes Supplement

This chapter supplements the standard user documentation, including the *Sun Cluster Data Services 5/03 Release Notes* shipped with the Sun™ Cluster 3.1 product. These “online release notes” provide the most current information on the Sun Cluster 3.1 product. This chapter includes the following information.

- “Revision Record” on page 13
- “New Features” on page 14
- “Restrictions and Requirements” on page 15
- “Known Problems” on page 15
- “Known Documentation Problems” on page 16

Revision Record

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

TABLE 2-1 Sun Cluster Data Services 3.1 5/03 Release Notes Supplement Revision Record: Year 2003

Revision Date	New Information
July 2003	Procedures to enable Sun Cluster Support for Oracle Parallel Server/Real Application Clusters on a subset of cluster nodes. See “Sun Cluster Support for Oracle Parallel Server/Real Application Clusters on a Subset of Cluster Nodes” on page 14. Revised support for Multiple Masters configuration of Sun Cluster HA for Sun ONE Application Server. See “Sun Cluster 3.1 Data Service for Sun ONE Application Server” on page 16.

New Features

In addition to features documented in *Sun Cluster Data Services 5/03 Release Notes*, this release now includes support for the following features.

Sun Cluster Support for Oracle Parallel Server/Real Application Clusters on a Subset of Cluster Nodes

You can enable Sun Cluster Support for Oracle Parallel Server/Real Application Clusters on a subset of cluster nodes. Install the data service packages only on the nodes that are configured to run Oracle Parallel Server/Real Application Clusters. You are not required to install data service packages on nodes that will not run Oracle Parallel Server/Real Application Clusters. For a list of packages and installation instructions, see *Sun Cluster 3.1 Data Service for Oracle Parallel Server/Real Application Clusters*.

Note – Restrictions apply when you use Sun Cluster Support for Oracle Parallel Server/Real Application Clusters with hardware RAID support or VxVM with the cluster feature. The Sun Cluster Support for Oracle Parallel Server/Real Application Clusters software must be installed only on the cluster nodes that are directly attached to the shared storage used by Oracle Parallel Server/Real Application Clusters.

Adding Sun Cluster Support for Oracle Parallel Server/Real Application Clusters to Selected Nodes

Add Sun Cluster Support for Oracle Parallel Server/Real Application Clusters to selected nodes in the following situations:

- You are adding nodes to a cluster and you plan to run Sun Cluster Support for Oracle Parallel Server/Real Application Clusters on the nodes.
- You are enabling Sun Cluster Support for Oracle Parallel Server/Real Application Clusters on a node.

To add Sun Cluster Support for Oracle Parallel Server/Real Application Clusters to selected nodes, the required data service software packages on the nodes. The storage management scheme that you are using determines which packages to install. For installation instructions, see *Sun Cluster 3.1 Data Service for Oracle Parallel Server/Real Application Clusters*.

Removing Sun Cluster Support for Oracle Parallel Server/Real Application Clusters From a Node

To remove Sun Cluster Support for Oracle Parallel Server/Real Application Clusters from selected nodes, remove software packages from the nodes. The storage management scheme that you are using determines which packages to remove.

▼ *How to Remove Sun Cluster Support for Oracle Parallel Server/Real Application Clusters From a Node*

1. **Become superuser.**
2. **Boot the nodes from which you are removing Sun Cluster Support for Oracle Parallel Server/Real Application Clusters in noncluster mode.**
3. **Uninstall from each node the Sun Cluster Support for Oracle Parallel Server/Real Application Clusters software packages for the storage management scheme that you are using.**
 - If you are using VxVM with the cluster feature, type the following command:
`# pkgrm SUNWscucm SUNWudlm SUNWudlmr SUNWcvmr SUNWcvm`
 - If you are using hardware RAID support, type the following command:
`# pkgrm SUNWscucm SUNWudlm SUNWudlmr SUNWschwr`
 - If you are using the cluster file system, type the following command:
`# pkgrm SUNWscucm SUNWudlm SUNWudlmr`

Restrictions and Requirements

The following restrictions and requirements have been added or updated since the Sun Cluster 3.1 release.

There are no restrictions and requirements at this time.

Known Problems

In addition to known problems documented in *Sun Cluster Data Services 5/03 Release Notes*, the following known problems affect the operation of the Sun Cluster 3.1 release.

There are no known problems at this time.

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 Data Services 5/03 Release Notes*.

Sun Cluster 3.1 Data Service for Sun ONE Application Server

Do not configure the Sun Cluster HA for Sun ONE Application Server as a resource that is mastered on multiple nodes at the same time. The multiple masters configuration is not supported. Only the failover configuration is supported. For information about supported configurations, contact your Sun Enterprise Services representative.

Sun Cluster 3.0 5/02 Release Notes Supplement

This document supplements the standard user documentation, including the *Sun Cluster 3.0 5/02 Release Notes* shipped with the Sun™ Cluster 3.0 product. These “online release notes” provide the most current information on the Sun Cluster 3.1 product. This document includes the following information.

- “Revision Record” on page 17
- “New Features” on page 25
- “Restrictions and Requirements” on page 28
- “Known Problems” on page 31
- “Known Documentation Problems” on page 34

Revision Record

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

TABLE 3-1 Sun Cluster 3.0 5/02 Release Notes Supplement Revision Record: Year 2003

Revision Date	New Information
July 2003	

TABLE 3-1 Sun Cluster 3.0 5/02 Release Notes Supplement Revision Record: Year 2003
(Continued)

Revision Date	New Information
June 2003	<p>Procedures to upgrade a Sun Cluster 3.0 configuration to Sun Cluster 3.1 software, including upgrading from Solaris 8 to Solaris 9 software. See Appendix F.</p> <p>Procedures to support Sun StorEdge 6320 storage systems. See “Installing and Maintaining a Sun StorEdge 6320 System” in <i>Sun Cluster 3.x With Sun StorEdge 6320 System Manual</i>.</p> <p>Sun StorEdge 6120 storage arrays in dual-controller configurations and Sun StorEdge 6320 storage systems are limited to four nodes and 16 LUNs. See “Bug ID 4840853” on page 11.</p> <p>Procedures to support Sun StorEdge 3510 FC storage device. See the <i>Sun Cluster 3.x With Sun StorEdge 3510 FC Array Manual</i> .</p> <p>Sun StorEdge 3510 FC storage arrays are limited to 256 LUNs per channel. See “Bug ID 4867584” on page 12.</p> <p>Sun StorEdge 3510 FC storage arrays are limited to one node per channel. See “Bug ID 4867560” on page 12.</p>
May 2003	<p>How to create node-specific files and directories for use with Oracle Parallel Server/Real Application Clusters on the cluster file system. See “Creating Node-Specific Files and Directories for Use With Oracle Parallel Server/Real Application Clusters Software on the Cluster File System” on page 47 for more information.</p> <p>New bge(7D) Ethernet adapter requires patches and modified installation procedure. See “BugID 4838619” on page 34 for more information.</p> <p>Increased stack-size settings are required when using VxFS. See “Bug ID 4662264” on page 31 for more information.</p>

TABLE 3-1 Sun Cluster 3.0 5/02 Release Notes Supplement Revision Record: Year 2003
(Continued)

Revision Date	New Information
April 2003	<p data-bbox="672 422 1422 499">Procedures to support Sun StorEdge 6120 storage arrays. See “Installing and Maintaining a Sun StorEdge 6120 Array” in <i>Sun Cluster 3.x With Sun StorEdge 6120 Array Manual</i>.</p> <p data-bbox="672 527 1422 575">Added VxVM Dynamic Multipathing (DMP) restrictions. See “Dynamic Multipathing (DMP)” on page 28 for more information.</p> <p data-bbox="672 602 1422 680">Support removed for the dual-bus, multi-initiator configuration of the Sun StorEdge 3310 JBOD storage array. See “BugId 4818874” on page 33 for more information.</p> <p data-bbox="672 707 1422 756">PCI Dual Ultra3 SCSI host adapter needs jumpers set for manual termination. See “BugId 4836405” on page 33 for more information.</p> <p data-bbox="672 783 1422 856">Added information on support for Oracle Parallel Server/Real Application Clusters on the cluster file system. See “Support for Oracle Parallel Server/Real Application Clusters on the Cluster File System” on page 26.</p> <p data-bbox="672 884 1422 982">Added information on using the Sun Cluster LogicalHostname resource with Oracle Parallel Server/Real Application Clusters. See “Using the Sun Cluster LogicalHostname Resource With Oracle Parallel Server/Real Application Clusters” on page 46.</p> <p data-bbox="672 1010 1422 1083">Sun Cluster HA for SAP now supports the SAP J2EE engine and SAP Web dispatcher configurations. For more information, see “Configuring an SAP J2EE Engine Cluster and an SAP Web Dispatcher” on page 44.</p> <p data-bbox="672 1110 1422 1163">Revised procedures on how to for install and configure Sun Cluster HA for SAP liveCache. See Appendix B.</p>
March 2003	<p data-bbox="672 1190 1422 1243">Revised support for installation of the Remote Shared Memory Reliable Datagram Transport (RSMRDT) driver. See Appendix D.</p> <p data-bbox="672 1270 1422 1344">Revised “How to Register and Configure Sun Cluster for SAP liveCache” procedure. See <i>Sun Cluster 3.1 Data Service Planning and Administration Guide</i>.</p> <p data-bbox="672 1371 1422 1423">Documentation bug in <code>sconf_transp_adap_sci(1M)</code> man page. See “<code>sconf_transp_adap_sci</code> Man Page” on page 53.</p> <p data-bbox="672 1451 1422 1514">Updated revised procedure on how to replace a disk drive in a StorEdge A5x00 storage array. See the <i>Sun Cluster 3.x With Sun StorEdge A5x00 Array Manual</i>.</p>

TABLE 3-1 Sun Cluster 3.0 5/02 Release Notes Supplement Revision Record: Year 2003
(Continued)

Revision Date	New Information
February 2003	<p>Revised procedures to support Sun Cluster HA for SAP on SAP 6.20. See Appendix E.</p> <p>Virtual Local Area Network (VLAN) support expanded. See the <i>Sun Cluster 3.x Hardware Administration Manual</i>.</p> <p>Procedures to support Sun StorEdge 9900 Dynamic Link Manager. See the <i>Sun Cluster 3.x With Sun StorEdge 9900 Series Storage Device Manual</i>.</p> <p>Revised <code>scconf_transp_adap_wrsm(1M)</code> man page to support a Sun Fire Link—based cluster interconnect. See “<code>scconf_transp_adap_wrsm Man Page</code>” on page 53.</p> <p>Procedures to support a Sun Fire Link—based cluster interconnect. See the <i>Sun Cluster 3.x Hardware Administration Manual</i>.</p>
January 2003	<p>Requirements on how to change connectivity to a quorum device. “Changing Quorum Device Connectivity” on page 11</p> <p>Support for daisy-chaining Sun StorEdge A1000 storage arrays. See the <i>Sun Cluster 3.x With Sun StorEdge A1000 or Netra st A1000 Array Manual</i>.</p> <p>Support for Cluster interconnects over Virtual Local Area Networks. See the <i>Sun Cluster 3.x Hardware Administration Manual</i>.</p> <p>Support for installation of the Remote Shared Memory Reliable Datagram Transport (RSMRDT) driver. See Appendix D.</p>

TABLE 3-2 Sun Cluster 3.x Release Notes Supplement Revision Record: Year 2002

Revision Date	New Information
December 2002	<p>Revised procedures on how to for install and configure Sun Cluster HA for SAP liveCache. See Appendix B.</p>

TABLE 3-2 Sun Cluster 3.x Release Notes Supplement Revision Record: Year 2002
(Continued)

Revision Date	New Information
November 2002	<p>Revised <code>SUNW.HASStoragePlus.5</code> man page to correct the Notes section and include <code>FilesystemCheckCommand</code> extension property. See “<code>SUNW.HASStoragePlus.5</code>” on page 54.</p> <p>Sun Cluster HA for Sun ONE Web Server now supports Sun ONE Proxy Server. See “Support for Sun ONE Proxy Server” on page 46.</p> <p>Name to use to configure SCI-PCI adapters for the cluster interconnect. See “Names for SCI-PCI Adapters” on page 42.</p> <p>Requirements for storage topologies. See “Storage Topologies Replaced by New Requirements” on page 29.</p> <p>Support for Dynamic Reconfiguration with the Sun Fire V880 system and Sun Cluster software. See <i>Sun Cluster 3.x Hardware Administration Manual</i>.</p> <p>Correction to the planning statement on how to connect quorum devices to nodes. See “Quorum Device Connection to Nodes” on page 43.</p> <p>Removal of the step on how to add nodes to the authentication list before you install VERITAS Volume Manager. See “New Features” on page 25.</p> <p>Package dependency to upgrade Sun Cluster HA for NFS from Sun Cluster 2.2 to Sun Cluster 3.0 software. See “Upgrading the Sun Cluster HA for NFS Data Service” on page 42.</p>
October 2002	<p><code>/etc/iu.ap</code> file to support the <code>ce</code> adapter. See “Installing and Maintaining Public Network Hardware” in <i>Sun Cluster 3.x Hardware Administration Manual</i>.</p> <p>Procedures to support Sun StorEdge 3310 RAID storage arrays. See “Installing and Maintaining Sun StorEdge 3310 SCSI RAID Array” in <i>Sun Cluster 3.x With Sun StorEdge 3310 Array Manual</i>.</p> <p>Procedures to support Sun StorEdge 3310 JBOD storage arrays. See “Installing and Maintaining Sun StorEdge 3310 SCSI JBOD Arrays” in <i>Sun Cluster 3.x With Sun StorEdge 3310 Array Manual</i>.</p> <p>Relaxed requirements for shared storage. See “Shared Storage Restriction Relaxed” on page 29.</p> <p>Procedure on how to replace a SCI-PCI host adapter in a running cluster. See <i>Sun Cluster 3.x Hardware Administration Manual</i>.</p> <p>Revised procedure on how to replace a disk drive in a non-RAID storage device. For the Sun Cluster documentation for the storage device, see <i>Sun Cluster 3.x Hardware Administration Collection</i>.</p> <p>Revised swap partition requirements. See “New Guidelines for the swap Partition” on page 26.</p>

TABLE 3-2 Sun Cluster 3.x Release Notes Supplement Revision Record: Year 2002
(Continued)

Revision Date	New Information
September 2002	<p>IP address configuration requirement for Sun Fire 15000 systems. See “IP Address Requirement for Sun Fire 15000 Systems” on page 43.</p> <p>Corrected cross-reference between uninstall procedures. See “How to Uninstall Sun Cluster Software From a Cluster Node (5/02)” on page 52.</p>
August 2002	<p>Restriction on EMC storage use in a two node configuration. See “EMC Storage Restriction” on page 29.</p>
July 2002	<p>Revised procedure to upgrade to the Sun Cluster 3.0 5/02 release from any previous version of Sun Cluster 3.0 software. See “How to Upgrade to the Sun Cluster 3.0 5/02 Software Update Release” on page 36.</p> <p>Revised procedure on how to replace a disk drive in StorEdge A5x00 storage array. See the <i>Sun Cluster 3.x With Sun StorEdge A5x00 Array Manual</i>.</p> <p>Requirements for ATM support with Sun Cluster 3.0 5/02. See “ATM with Sun Cluster 3.0 5/02” on page 34</p> <p>Sun Cluster Security Hardening support for Solaris 9. See “Security Hardening for Solaris 9” on page 25.</p>

TABLE 3-2 Sun Cluster 3.x Release Notes Supplement Revision Record: Year 2002
(Continued)

Revision Date	New Information
June 2002	<p>Restriction on concurrent upgrade of Solaris 9 and Sun Cluster 3.0 5/02 software. See “Framework Restrictions and Requirements” on page 30.</p> <p>Revised appendix to support Sun StorEdge 9970 system and Sun StorEdge 9980 system with Sun Cluster software. See the <i>Sun Cluster 3.x With Sun StorEdge 9900 Series Storage Device Manual</i>.</p> <p>Procedures to support Sun StorEdge D2 storage systems. See <i>Sun Cluster 3.x With Sun StorEdge D2 Array Manual</i>.</p> <p>Revised procedures to support Sun StorEdge T3/T3+ Partner Group and Sun StorEdge 3900 storage arrays in a 4-node configuration. See “Sun StorEdge T3/T3+ Partner Group and Sun StorEdge 3900 Storage Devices Supported in a Scalable Topology.” on page 36.</p> <p>Updated procedures to support Sun Cluster software on Sybase 12.0 64-bit version. See Appendix C.</p> <p>Documentation bug in the Sun Cluster Hardware Guide. See “Failover File System (HAStoragePlus)” on page 25.</p> <p>Documentation bug in the Sun Cluster Hardware Guide. See “Changing Quorum Device Connectivity” on page 29.</p> <p>Documentation bug in the Sun Cluster Hardware Guide: ce Sun Ethernet Driver Considerations. See “Installing and Maintaining Public Network Hardware” in <i>Sun Cluster 3.x Hardware Administration Manual</i>.</p> <p>Documentation bug in the Sun Cluster Hardware Guide: Hard zone configuration changed. See the <i>Sun Cluster 3.x With Sun StorEdge 3900 or 6900 Series System Manual</i>.</p> <p>Updated procedures to support Apache version 2.0. See “Apache 2.0” on page 26.</p>

TABLE 3-2 Sun Cluster 3.x Release Notes Supplement Revision Record: Year 2002
(Continued)

Revision Date	New Information
May 2002	<p>Oracle UDLM requirement. See "Oracle UDLM Requirement" on page 30.</p> <p>Restriction on IPv6 and IP Network Multipathing. See "Framework Restrictions and Requirements" on page 30.</p> <p>Failover File System (HAStoragePlus) support. See "Quorum Device Connection to Nodes" on page 43.</p> <p>RAID-5 support on Sun StorEdge 99x0 storage arrays. See "RAID 5 on Sun StorEdge 99x0 Storage Arrays" on page 26.</p> <p>Correction to BugId 4662264 workaround. See "BugId 4662264" on page 53.</p> <p>"Bug ID 4346123" on page 31, cluster file system might not mount after multiple failures.</p> <p>"Bug ID 4665886" on page 32, mapping a file into the address space with <code>mmap(2)</code> and then issuing a <code>write(2)</code> call to the same file results in a recursive mutex panic.</p> <p>"Bug ID 4668496" on page 32, Solaris™ Volume Manager replicas need more space.</p> <p>"Bug ID 4680862" on page 32, node needs access to the highly available local file system managed by HAStoragePlus.</p> <p>Documentation bug in the Sun Cluster data services collection. See "Configuring Sun ONE Web Server" on page 45.</p> <p>Space requirement for Solaris Volume Manager. See "Solaris Volume Manager Replica Space Requirement" on page 43.</p> <p>Information and procedures on how to use the new scalable cluster topology. See Appendix A.</p> <p>Documentation bug in the Sun Cluster Hardware Guide. See the <i>Sun Cluster 3.x Hardware Administration Manual</i>.</p>

TABLE 3-2 Sun Cluster 3.x Release Notes Supplement Revision Record: Year 2002
(Continued)

Revision Date	New Information
	Documented campus clustering configuration information to include support for the Sun StorEdge 9910 storage device and Sun StorEdge 9960 storage device. See the <i>Sun Cluster 3.x With Sun StorEdge 9900 Series Storage Device Manual</i> .

New Features

In addition to features documented in *Sun Cluster 3.0 5/02 Release Notes*, this release now includes support for the following features.

Security Hardening for Solaris 9

Sun Cluster Security Hardening now supports data services in a Solaris 8 and Solaris 9 environment. The Sun Cluster Security Hardening documentation is available at <http://www.sun.com/security/blueprints>. From this URL, scroll down to the Architecture heading to locate the article "Securing the Sun Cluster 3.0 Software."

Failover File System (HAStoragePlus)

Failover File System (HAStoragePlus) is now supported in Sun Cluster 3.0 5/02 release. See the *Sun Cluster 3.0 5/02 Supplement* for information about this new feature.

The `FilesystemMountPoints` extension property can be used to specify a list of one or more file system mount points. This list can consist of both local and global file system mount points.

RAID 5 on Sun StorEdge 99x0 Storage Arrays

RAID level 5 is supported on Sun StorEdge 99x0 storage arrays with multipathing and Sun adapters.

Apache 2.0

Sun Cluster 3.1 now supports Apache version 2.0. For Apache version 2.0, the procedure for configuring the `httpd.conf` configuration file has changed as follows. (See the Sun Cluster data services collection for the complete procedure.)

- The `ServerName` directive specifies the hostname and the port.
- The `BindAddress` and `Port` directives have been replaced with the `Listen` directive. The `Listen` directive must use the address of the logical host or shared address.
- The `Servertype` directive no longer exists.

New Guidelines for the swap Partition

The amount of swap space allocated for Solaris and Sun Cluster software combined must be no less than 750 Mbytes. For best results, add at least 512 Mbytes for Sun Cluster software to the amount required by the Solaris operating environment. In addition, allocate additional swap space for any third-party applications you install on the node that also have swap requirements. See your third-party application documentation for any swap requirements.

Support for Oracle Parallel Server/Real Application Clusters on the Cluster File System

You can use Oracle Parallel Server/Real Application Clusters with the cluster file system.

Pre-Installation Considerations

Oracle Parallel Server/Real Application Clusters is a scalable application that can run on more than one node concurrently. You can store all of the files that are associated with this application on the cluster file system, namely:

- Binary files
- Control files
- Data files

- Log files
- Configuration files

For optimum I/O performance during the writing of redo logs, ensure that the following items are located on the same node:

- The Oracle Parallel Server/Real Application Clusters database instance
- The primary of the device group that contains the cluster file system that holds the following logs of the database instance:
 - Online redo logs
 - Archived redo logs

For other pre-installation considerations that apply to Sun Cluster Support for Oracle Parallel Server/Real Application Clusters, see “Overview of the Installation and Configuration Process” in *Sun Cluster 3.0 12/01 Data Services Installation and Configuration Guide*.

How to Use the Cluster File System

To use the cluster file system with Oracle Parallel Server/Real Application Clusters, create and mount the cluster file system as explained in “Configuring the Cluster” in *Sun Cluster 3.0 12/01 Software Installation Guide*. When you add an entry to the `/etc/vfstab` file for the mount point, set UNIX file system (UFS) file system specific options for various types of Oracle files as shown in the following table.

TABLE 3-3 UFS File System Specific Options for Oracle Files

File Type	Options
RDBMS data files, log files, control files	global, logging, forcedirectio
Oracle binary files, configuration files	global, logging

▼ How to Install Sun Cluster Support for Oracle Parallel Server/Real Application Clusters Packages With the Cluster File System

To complete this procedure, you need the Sun Cluster 3.1 Agents CD-ROM. Perform this procedure on all of the cluster nodes that can run Sun Cluster Support for Oracle Parallel Server/Real Application Clusters.

Note – Due to the preparation that is required prior to installation, the `scinstall(1M)` utility does not support automatic installation of the data service packages.

1. Load the Sun Cluster 3.1 Agents CD-ROM into the CD-ROM drive.
2. Become superuser.
3. On all of the nodes, run the following command to install the data service packages.

```
# pkgadd -d \  
  /cdrom/scdataservices_3_1_u3/components/  
  SunCluster_Oracle_Parallel_Server_3.0_u3/Packages \  
  SUNWscucm SUNWudlm SUNWudlmr
```



Caution – Before you reboot the nodes, you must ensure that you have correctly installed and configured the Oracle UDLM software. For more information, see “Installing the Oracle Software” in *Sun Cluster 3.0 12/01 Data Services Installation and Configuration Guide*.

Where to Go From Here

Go to “Installing the Oracle Software” in *Sun Cluster 3.0 12/01 Data Services Installation and Configuration Guide* to install the Oracle UDLM and Oracle RDBMS software.

Restrictions and Requirements

The following restrictions and requirements have been added or updated since the Sun Cluster 3.0 12/01 release.

Dynamic Multipathing (DMP)

Use of VxVM Dynamic Multipathing (DMP) with Sun Cluster 3.1 software to manage multiple paths from the same node is not supported. From VxVM 3.2 onward it is no longer possible to disable the installation of DMP. But having it in the I/O stack on systems with only a single path per node poses no problems. However, if you use VxVM in a configuration with multiple paths per node, then you must use another multipathing solution, such as MPxIO or EMC PowerPath.

Changing Quorum Device Connectivity

When you increase or decrease the number of node attachments to a quorum device, the quorum vote count is not automatically recalculated. You can reestablish the correct quorum vote if you remove all quorum devices and then add them back into the configuration.

Storage Topologies Replaced by New Requirements

Sun Cluster 3.1 software now supports open topologies. You are no longer limited to the storage topologies listed in the Sun Cluster concepts documentation document.

Use the following guidelines to configure your cluster.

- Sun Cluster supports a maximum of eight nodes in a cluster, regardless of the storage configurations that you implement.
- A shared storage device can connect to as many nodes as the storage device supports.
- Shared storage devices do not need to connect to all nodes of the cluster. However, these storage devices must connect to at least two nodes.

Shared Storage Restriction Relaxed

Sun Cluster 3.1 now supports greater than three-node cluster configurations without shared storage devices. Two-node clusters are still required to have a shared storage device to maintain quorum. This storage device does not need to perform any other function.

EMC Storage Restriction

The quorum device access mode is not automatically set to `scsi-3` in the following situations:

- After applying the core patch 110648-20 or later in a two node cluster with an EMC Powerpath configured quorum disk.
- After upgrading from Sun Cluster 3.0 12/01 software to Sun Cluster 3.0 05/02 software in a two node cluster with an EMC Powerpath configured quorum disk.

Note – This is a problem only for a multipath quorum device configured with EMC Powerpath in a two node configuration. The problem is characterized by a value of NULL being printed for the quorum device access mode property.

To fix the property setting after applying the patch or performing the upgrade, use the `scsetup` command to remove the existing quorum disk and add it back to the configuration. Removing and adding back the quorum disk will correct the Sun Cluster software to use `scsi-3` PGR for reserving quorum disks. To verify that the quorum device access mode is set correctly, run `scconf -p` to print the configuration.

Framework Restrictions and Requirements

- **Upgrade to Solaris 9** – Upgrade to Solaris 9 software during upgrade to Sun Cluster 3.0 5/02 software is not supported. You can only upgrade to subsequent, compatible versions of the Solaris 8 operating environment during upgrade to Sun Cluster 3.0 5/02 software. To run Sun Cluster 3.0 5/02 software on the Solaris 9 operating environment, you must perform a new installation of the Solaris 9 version of Sun Cluster 3.0 5/02 software after the nodes are upgraded to Solaris 9 software.
- **IPv6** – This is not supported.
- **IP Network Multipathing** – This is not supported.

Oracle UDLM Requirement

Oracle RAC running on Sun Cluster software requires that the Oracle UDLM be at least version 3.3.4.5, which ships with the Oracle 9.2 release.



Caution – If you do not have this revision or higher, you might encounter a problem during a cluster reconfiguration where the reconfiguration process will hang, leaving all nodes in the cluster unable to provide Oracle RAC database service. You can fix this problem by ensuring that your Oracle UDLM is at least version 3.3.4.5. This problem and fix are documented in Oracle Bug #2273410.

You can determine the version of Oracle UDLM currently installed on your system by running the following command.

```
pkginfo -l ORCLudlm | grep VERSION
```

The version of the Oracle UDLM currently installed on your system also appears in the following file.

```
/var/opt/SUNWcluster/dlm_node_name/logs/dlmstart.log
```

The version information appears just before the Copyright (c) line. Look for the latest occurrence of this information in the file. If you do not have this version of the Oracle UDLM package, please contact Oracle Support to obtain the latest version.

Known Problems

In addition to known problems documented in *Sun Cluster 3.0 5/02 Release Notes*, the following known problems affect the operation of the Sun Cluster 3.0 12/01 release.

Bug ID 4346123

Problem Summary: When booting a cluster node after multiple failures, a cluster file system might fail to mount automatically from its `/etc/vfstab` entry, and the boot process will place the node in an administrative shell. Running the `fsck` command on the device might yield the following error.

```
Can't roll the log for /dev/global/rdisk/dXsY
```

Workaround: This problem might occur when the global device is associated with a stale cluster file system mount. Run the following command, and check if the file system shows up in an error state to confirm a stale mount.

```
# /usr/bin/df -k
```

If the global device is associated with a stale cluster file system mount, unmount the global device. If any users of the file system exist on any of the nodes, the unmount cannot succeed. Run the following command on each node to identify current users of the file system.

```
# /usr/sbin/fuser -c mountpoint
```

If there are users of the file system, terminate those users' connection to the file system. Run the `share(1M)` command to confirm that the file system is not NFS- shared by any node.

Bug ID 4662264

Problem Summary: To avoid panics when using VxFS with Sun Cluster software, the default thread stack size must be greater than the VxFS default value of 0x4000.

Workaround: Increase the stack size by putting the following lines in the `/etc/system` file:

```
set rpcmod:svc_default_stksize=0x8000
set lwp_default_stksize=0x6000
```

After installing VxFS packages, verify that VxFS installation has not added similar statements to the `/etc/system` file. If multiple entries exist, resolve them to one statement per variable, using these higher values.

Bug ID 4665886

Problem Summary: Mapping a file into the address space with `mmap(2)` and then issuing a `write(2)` call to the same file results in a recursive mutex panic. This problem was identified in a cluster configuration running the iPlanet™ Mail Server.

Workaround: There is no workaround.

Bug ID 4668496

Problem Summary: The default `JumpStart™ profile` file allocates 10 Mbytes to slice 7. If you use Solaris 9 software with Solstice DiskSuite, this amount of space is not enough for Solstice DiskSuite replicas. Solaris 9 software with Solstice DiskSuite requires at least 20 Mbytes.

Workaround: Edit the default `profile` file to configure slice 7 of the system disk with 20 Mbytes of space, instead of 10 Mbytes. This workaround is only necessary if you install Solaris 9 software with Solstice DiskSuite

Bug ID 4680862

Problem Summary: When you install Oracle or Sybase binaries and configuration files on a highly available local file system managed by HAStoragePlus, the node that does not have access to this file system fails validation. The result is that you cannot create the resource.

Workaround: Create a symbolic link named `/usr/cluster/lib/hasp_check` to link to the `/usr/cluster/lib/scdsbuilder/src/scripts/hasp_check` file.

Bug ID 4779686

Problem Summary: Availability Suite 3.1 does not support the Sun Cluster 3.0 HAStoragePlus resource.

Workaround: If you intend to implement Availability Suite 3.1 and failover file system, use an HAStorage resource in the light-weight resource group that includes the Availability Suite logical host. For the application resource group, use HAStoragePlus. This allows you to use a failover file system for application performance and also use Availability Suite 3.1 to back up the disk blocks under the failover file system.

BugId 4818874

Problem Summary: When used in a clustered environment, the Sun StorEdge 3310 JBOD array relies on the cluster nodes to provide SCSI bus termination. As termination power is not supplied from the array's "IN" ports, if the server connected to these loses power bus termination will be lost. This in turn may result in the remaining cluster node losing access to the shared storage on that bus.

Workaround: The StorEdge 3310 JBOD array is not supported in a split-bus configuration. To fix affected configurations, add an additional StorEdge 3310 JBOD array to the cluster configuration.

BugId 4836405

Problem Summary: When using the PCI Dual Ultra3 SCSI host adapter in a clustered environment, the host adapter jumpers for each port must be set for manual SCSI termination. If the ports are not set to manual SCSI termination, a loss of power to one host could prevent correct SCSI bus operation and might result in loss of access to all SCSI devices attached to that bus from the remaining host.

Workaround: When using the PCI Dual Ultra3 SCSI host adapter in a clustered environment, set the jumpers on the host adapter to manual SCSI termination. This setting causes the host adapter to activate its built-in SCSI terminators, whether or not the host adapter receives PCI bus power.

The jumper settings needed for manual termination are listed below.

- SCSI bus 2 (external SCSI connector nearest to the PCI slot)
 - J4: 2-3 (factory default 2-3)
 - J5: 2-3 (factory default 2-3)
- SCSI bus 1 (internal SCSI connector and external SCSI connector furthest from the PCI slot)
 - J8: 2-3 (factory default 1-2)
 - J9: 2-3 (factory default 1-2)

See the host adapter documentation for further information.

BugID 4838619

Problem Summary: Without a patch, Sun Cluster will not recognize bge(7D) Ethernet adapters.

Workaround: If you plan to use bge(7D) Ethernet adapters as cluster interconnects in your Sun Cluster configuration, you will need to install patches and use a modified installation procedure. The onboard Ethernet ports on the Sun Fire V210 and V240 are examples of bge(7D) Ethernet adapters.

If you use Solaris 8, install the following patches.

- 110648-28 or later (Sun Cluster 3.0: Core/Sys Admin)
- 112108-07 or later (Required for SunPlex Manager use)

If you use Solaris 9, install the following patches.

- 112563-10 or later (Sun Cluster 3.0: Core/Sys Admin)
- 114189-01 or later (Required for SunPlex Manager use)

For the modified installation procedure, refer to the patch's README file.

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 5/02 Release Notes*.

Hardware Guide

The following subsections describe omissions or new information that will be added to the next publishing of the Hardware Guide.

ATM with Sun Cluster 3.0 5/02

ATM is supported with Sun Cluster 3.0 5/02 as a public network interface to be used in LAN Emulation (LANE) mode only. Use the SunATM 5.0 version to run on Solaris 8.

Use the following network, ATM card, and LANE instance guidelines to configure ATM with Sun Cluster 3.0 5/02. For additional configuration information, see the *Platform Notes: The SunATM Driver Software*.

Network Configuration Guidelines

In order to support ATM LANE on Sun Cluster, an ATM capable router and switch are required. The router needs to provide LANE services, with 1 ELAN for each set of nodes. Configure the router to respond to ALLROUTERS (224.0.0.2) and the ALLHOSTS (224.0.0.1) pings. The ATM switch should have PNNI (Private Network-Node Interface) enabled.

The router provides Emulated LAN (ELAN) service to the cluster nodes and clients. The clients can belong to a different ELANs but the cluster nodes must be part of the same ELAN.

ATM Card Guidelines

To use ATM as a public network adapter, Sun Cluster requires at least one ATM card per NAFO group. For high availability, you can eliminate the potential single point of failure by using more than one ATM card per NAFO group.

LANE Instance Configuration Guidelines

Perform the following tasks to configure LANE instances.

- Create one LANE instance on each ATM card.
- All LANE instances in a NAFO group must be configured on the same ELAN. For example, all LANE instances in NAFO1 must be in the same ELAN on all cluster nodes.
- Configure the primary LANE interface using the `/etc/hostname.lanen` file. This file is necessary, but will cause warning messages to display at boot up on SunATM 5.0. The following example is of the console messages. These messages can be ignored.

```
Rebooting with command: boot
Boot device: diskbrd:a File and args:
SunOS Release 5.8 Version Generic_108528-13 64-bit
Copyright 1983-2001 Sun Microsystems, Inc. All rights reserved.
ip_rput_dlpi(lan1): DL_ERROR_ACK for DL_ATTACH_REQ(11), errno 8, unix 0
ip_rput_dlpi(lan1): DL_ERROR_ACK for DL_BIND_REQ(1), errno 3, unix 71
ip_rput_dlpi(lan1): DL_ERROR_ACK for DL_PHYS_ADDR_REQ(49), errno 3, unix 71
ip_rput_dlpi(lan1): DL_ERROR_ACK for DL_UNBIND_REQ(2), errno 3, unix 71
ip_rput_dlpi(lan1): DL_ERROR_ACK for DL_DETACH_REQ(12), errno 3, unix 71
ifconfig: SIOCSLIFNAME for ip: lan1: Protocol error
moving addresses from failed IPv4 interfaces: lan1 (couldn't move, no
alternative interface).
Hostname: atm10
```

- Assign an IP address to the primary LANE interface in the `atmconfig` file.

Note – Do not assign an IP address to the secondary, backup LANE interface.

The following example shows an `atmconfig` file with the primary and secondary LANE interfaces configured. Note the IP address is assigned only to the primary LANE interface.

```
ba0      3.1      -      -  
ba0      SONET    -      -      -  
ba0      -      -      1      atm20  
ba1      3.1      -      -  
ba1      SONET    -      -      -  
ba1      -      -      2      -
```

Sun StorEdge T3/T3+ Partner Group and Sun StorEdge 3900 Storage Devices Supported in a Scalable Topology.

The Sun StorEdge T3/T3+ Partner Group and Sun StorEdge 3900 storage devices are supported with 4-node connectivity in a cluster environment.

To configure and maintain these storage devices with 4-node connectivity, use the procedures listed in the storage device's chapter and repeat the steps for Node B on each additional node that connects to the storage device.

For the following node-related procedures, see Appendix A.

- "Adding a Cluster Node" on page 58
- "Removing a Cluster Node" on page 58
- "How to Remove Connectivity Between an Array and a Single Node, in a Cluster With Greater Than Two-Node Connectivity" on page 60

Software Installation Guide

The following subsections describe omissions or new information that will be added to the next publishing of the Software Installation Guide.

▼ How to Upgrade to the Sun Cluster 3.0 5/02 Software Update Release

Use the following procedure to upgrade any previous release of Sun Cluster 3.0 software to the Sun Cluster 3.0 5/02 update release.

Note – Do not use any new features of the update release, install new data services, or issue any administrative configuration commands until all nodes of the cluster are successfully upgraded.

1. Back up the shared data from all device groups within the cluster.

2. Get any necessary patches for your cluster configuration.

In addition to Sun Cluster software patches, get any patches for your hardware, Solaris operating environment, volume manager, applications, and any other software products currently running on your cluster. See the Sun Cluster release notes documentation for the location of Sun patches and installation instructions. You will apply the patches in different steps of this procedure.

3. From any node, view the current status of the cluster to verify that the cluster is running normally.

```
% scstat
```

See the `scstat(1M)` man page for more information.

4. Become superuser on one node of the cluster.

Upgrade only one node at a time.

5. Evacuate all resource groups and device groups that are running on the node to upgrade.

Specify the node that you are upgrading in the `node` argument of the following `scswitch` command:

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

`-s` Evacuates all resource groups and device groups

`-h node` Specifies the name of the node from which to evacuate resource groups and device groups (the node you are upgrading)

See the `scswitch(1M)` man page for more information.

6. Verify that the evacuation completed successfully.

```
# scstat -g -D
```

Ensure that the node you are upgrading is no longer the primary for any resource groups or device groups in the cluster.

7. Reboot the node into non-cluster mode.

Include the double dashes (`--`) in the command.

```
# reboot -- -x
```

8. Back up the system disk.

9. Determine whether any of the Cool Stuff CD packages are installed on the node.

To display the version of an installed package, use the following command:

```
# pkginfo -l package
```

The following table lists the packages from the Sun Cluster 3.0 GA Cool Stuff CD-ROM:

TABLE 3-4 Sun Cluster 3.0 GA Cool-Stuff Packages to Remove

Package	Version	Description
SUNWscrtw	3.0.0/2000.10.17.22.22	Resource Type Wizard
SUNWscsdk	3.0.0/2000.10.10.13.06	Data Service Software Development Kit
SUNWscset	3.0.0/2000.10.17.22.22	rgmsetup
SUNWscvxi	3.0.0/2000.10.17.22.22	Cluster VxVM setup

Remove any Cool Stuff CD-ROM packages found on the node. These packages will be replaced with supported versions in Sun Cluster 3.0 5/02 software.

```
# pkgrm package
```

10. Do you intend to upgrade Solaris 8 software?

Note – The cluster must already run on, or be upgraded to, at least the minimum required level of the Solaris 8 operating environment to support Sun Cluster 3.0 5/02 software. See the Info Documents page for Sun Cluster 3.0 software on <http://sunsolve.sun.com> for the latest Solaris support information.

- If yes, go to Step 11.
- If no, go to Step 12.

11. Upgrade Solaris 8 software.

a. Determine whether the following links already exist, and if so, whether the file names contain an uppercase K or S.

```
/etc/rc0.d/K16apache  
/etc/rc1.d/K16apache  
/etc/rc2.d/K16apache  
/etc/rc3.d/S50apache  
/etc/rcS.d/K16apache
```

If these links already exist **and** contain an uppercase K or S in the file name, no further action is necessary concerning these links. If these links do not exist, or if these links exist but contain a lowercase k or s in the file name, you will move aside these links in Step g.

b. Are you using the Maintenance Update upgrade method?

- If yes, skip to Step c.
- If no, temporarily comment out all global device entries in the `/etc/vfstab` file.

Do this to prevent the Solaris upgrade from attempting to mount the global devices. To identify global device entries, look for entries that contain `global` in the `mount-options` list.

c. Shut down the node to upgrade.

```
# shutdown -y -g0
ok
```

d. Follow instructions in the installation guide for the Solaris 8 update version you want to upgrade to.

Note – To reboot the node during Solaris software upgrade, always add the `-x` option to the command. This ensures that the node reboots into non-cluster mode. The following two commands boot a node into single-user non-cluster mode:

```
# reboot -- -sx
ok boot -sx
```

Do not reboot the node into cluster mode during or after Solaris software upgrade.

e. Are you using the Maintenance Update upgrade method?

- If yes, skip to Step f.
- If no, uncomment all global device entries that you commented out in the `/a/etc/vfstab` file.

f. Install any Solaris software patches and hardware-related patches, and download any needed firmware contained in the hardware patches.

Do not reboot yet if any patches require rebooting.

g. If the Apache links in Step a did not already exist or they contained a lowercase k or s in the file names before you upgraded Solaris software, move aside the restored Apache links.

Use the following commands to rename the files with a lowercase k or s:

```
# mv /a/etc/rc0.d/K16apache /a/etc/rc0.d/k16apache
# mv /a/etc/rc1.d/K16apache /a/etc/rc1.d/k16apache
# mv /a/etc/rc2.d/K16apache /a/etc/rc2.d/k16apache
# mv /a/etc/rc3.d/S50apache /a/etc/rc3.d/s50apache
# mv /a/etc/rcS.d/K16apache /a/etc/rcS.d/k16apache
```

Note – For the Maintenance Update upgrade method, the paths to the files do not begin with /a.

h. Reboot the node into non-cluster mode.

Include the double dashes (--) in the command.

```
# reboot -- -x
```

12. Determine whether the following packages are installed on the node.

```
# pkginfo SUNWscva SUNWscvr SUNWscvw SUNWscgds
```

Sun Cluster software upgrade requires that these packages exist on the node before upgrade begins. If any of these packages are missing, install them from the Sun Cluster 3.1 CD-ROM.

```
# cd /cdrom/suncluster_3_1/SunCluster_3.0/Packages
# pkgadd -d . SUNWscva SUNWscvr SUNWscvw SUNWscgds
```

13. Do you intend to use SunPlex Manager?

- If no, go to Step 14.
- If yes, ensure that the required Apache software packages are installed on the node.

```
# pkginfo SUNWapchr SUNWapchu
```

If any Apache software packages are missing, install them on the node from the Solaris CD-ROM.

```
# pkgadd -d . SUNWapchr SUNWapchu
```

14. Upgrade to the Sun Cluster 3.0 5/02 update software.

a. Insert the Sun Cluster 3.1 CD-ROM into the CD-ROM drive on the node.

If the volume daemon vold(1M) is running and configured to manage CD-ROM devices, it automatically mounts the CD-ROM on the /cdrom/suncluster_3_1 directory.

b. Change to the Tools directory.

```
# cd /cdrom/suncluster_3_1/SunCluster_3.0/Tools
```

c. Install the Sun Cluster 3.0 5/02 update patches.

```
# ./scinstall -u update
```

See the scinstall(1M) man page for more information.

d. Change to the CD-ROM root directory and eject the CD-ROM.

e. Install any Sun Cluster software patches.

- f. Verify that each Sun Cluster 3.0 5/02 update patch is installed correctly.
View the upgrade log file referenced at the end of the upgrade output messages.

15. Reboot the node into the cluster.

```
# reboot
```

16. Verify the status of the cluster configuration.

```
% scstat
```

17. Repeat Step 4 through Step 16 on each remaining cluster node, one node at a time.

18. Take offline all resource groups for the data services you will upgrade.

```
# scswitch -F -g resource-grp
-F                Take offline
-g resource-grp   Specifies the name of the resource group to take offline
```

19. Upgrade applications as needed.

Follow the instructions provided in your third-party documentation.

20. On each cluster node on which data services are installed, upgrade to the Sun Cluster 3.0 5/02 data services update software.

- a. Insert the Sun Cluster 3.1 Agents CD-ROM into the CD-ROM drive on the node.

- b. Install the Sun Cluster 3.0 5/02 data services update patches.

Use one of the following methods:

- To upgrade one or more specified data services, type the following command:

```
# scinstall -u update -s svc[,svc,...] -d cdrom-image
```

- To upgrade all data services present on the node, type the following command:

```
# scinstall -u update -s all -d cdrom-image
```

Note – The `-s all` option assumes that updates for all installed data services exist on the update release. If an update for a particular data service does not exist in the update release, that data service is not upgraded.

- c. Eject the CD-ROM.

- d. Install any Sun Cluster data service software patches.

- e. **Verify that each data service update patch is installed successfully.**

View the upgrade log file referenced at the end of the upgrade output messages.

- 21. After all data services on all cluster nodes are upgraded, bring back online the resource groups for each upgraded data service.**

```
# scswitch -Z -g resource-grp
-Z          Bring online
```

- 22. From any node, verify the status of the cluster configuration.**

```
% scstat
```

- 23. Restart any applications.**

Follow the instructions provided in your application's documentation.

Upgrading the Sun Cluster HA for NFS Data Service

In the procedure "How to Finish Upgrading Cluster Software" from the section "Upgrading From Sun Cluster 2.2 to Sun Cluster 3.0 Software," the following command upgrades the Sun Cluster HA for NFS data service:

```
# scinstall -u finish -q globaldev=DIDname \  
-d /cdrom/scdataservices_3_0_u3 -s nfs
```

This command requires that the `SUNWscnfs` package is already installed from the Sun Cluster 3.1 Agents CD-ROM on all nodes before you invoke the `scinstall` command. To ensure successful upgrade of the Sun Cluster HA for NFS data service, do the following:

- Ensure that the `SUNWscnfs` package is installed on all nodes of the cluster before you run this `scinstall` command.
- If the `scinstall` command fails because the `SUNWscnfs` package is missing from a node, install the `SUNWscnfs` package on all nodes from the Sun Cluster 3.1 Agents CD-ROM, then rerun the `scinstall` command.

Names for SCI-PCI Adapters

To configure SCI-PCI adapters for the cluster interconnect, specify `sciN` as the adapter name, for example, `sci0`. Do not use `scidN` as the adapter name.

Solaris Volume Manager Replica Space Requirement

Problem Summary: The *Sun Cluster 3.0 12/01 Software Installation Guide* tells you to set aside at least 10 Mbytes in slice 7 to use to create three Solaris Volume Manager replicas in that slice. However, Solaris Volume Manager replicas in Solaris 9 software require substantially more space than the 10 Mbytes required for Solstice DiskSuite replicas in Solaris 8 software.

Workaround: When you install Solaris 9 software, allocate at least 20 Mbytes to slice 7 of the root disk to accommodate the larger Solaris Volume Manager replicas.

IP Address Requirement for Sun Fire 15000 Systems

Before you install Sun Cluster software on a Sun Fire 15000 system, you must add the IP address of each domain console network interface to the `/etc/inet/hosts` file on each node in the cluster. Perform this task regardless of whether you use a naming service.

Quorum Device Connection to Nodes

In the Planning chapter, the following statement about quorum devices is incorrect:

Connection - Do not connect a quorum device to more than two nodes.

The statement should instead read as follows:

Connection – You must connect a quorum device to at least two nodes.

Node Authentication When Installing VERITAS Volume Manager

In the procedures “How to Install VERITAS Volume Manager Software and Encapsulate the Root Disk” and “How to Install VERITAS Volume Manager Software Only,” it is no longer necessary to first add cluster node names to the authentication list. You can therefore skip Step 3, “Add all nodes in the cluster to the cluster node authentication list.”

Data Services Installation and Configuration Guide

The following subsections describe omissions or new information that will be added to the next publishing of the *Data Service Installation and Configuration Guide*.

Configuring an SAP J2EE Engine Cluster and an SAP Web Dispatcher

Sun Cluster now supports the SAP J2EE engine cluster and SAP Web dispatcher components on the Sun Cluster environment. To use these components you must complete additional steps during your Sun Cluster HA for SAP installation and configuration.

- To configure a J2EE engine cluster with your Sun Cluster HA for SAP with a Central Instance, see “How to Configure an SAP J2EE Engine with your Sun Cluster HA for SAP with Central Instance” on page 44.
- To configure a J2EE engine cluster with your Sun cluster HA for SAP with an SAP Application Server, see “How to Configure an SAP J2EE Engine Cluster with your Sun Cluster HA for SAP with an Application Server” on page 45.
- To configure SAP Web dispatcher with your Sun Cluster HA for SAP agent, see “How to Configure a SAP Web Dispatcher with your Sun Cluster HA for SAP” on page 45.

The SAP J2EE engine is started by the SAP dispatcher which is under the protection of the Sun Cluster HA for SAP. If the SAP J2EE engine goes down, the SAP dispatcher will restart it.

The SAP Web dispatcher has the capability of auto restart. If the SAP Web dispatcher goes down, the SAP Web dispatcher watch dog process will restart. Currently, there is no Sun Cluster agent available for the SAP Web dispatcher.

▼ *How to Configure an SAP J2EE Engine with your Sun Cluster HA for SAP with Central Instance*

After you have completed the *How to Enable Failover SAP Instances to Run in a Sun Cluster* procedure in the *Sun Cluster HA for SAP* document, perform the following steps.

1. **Using the SAP J2EE Admintool GUI, change the ClusterHosts parameter to list all logical hosts for the application server and port pair under dispatcher/Manager/ClusterManager. For example,**

```
as1-1h:port;as2-1h:port ...
```

2. **Change the file `j2ee-install-dir/additionalproperties` as follows:**

```
com.sap.instanceId = logical-host-ci_SID_SYSNR
```

3. **Change the file `j2ee-install-dir/server/services/security/work/R3Security.properties` as follows:**

```
sapbasis.ashost = logical-host-ci
```

4. Change the file `SDM-dir/program/config/flow.xml`

```
host = logical-host-ci
```

▼ *How to Configure an SAP J2EE Engine Cluster with your Sun Cluster HA for SAP with an Application Server*

After you have completed the *How to Enable Failover SAP Instances to Run in a Sun Cluster* or *How to Install an SAP Scalable Application Server* procedure in the *Sun Cluster HA for SAP* document, perform the following steps.

1. Using the SAP J2EE Admintool GUI, change ClusterHosts parameter to list the logical host for the central instance and port pair under the dispatcher/Manager/ClusterManager.

```
logical-host-ci:port
```

2. Change the file `j2ee-install-dir/additionalproperties` as follows:

```
com.sap.instanceId = logical-host-as_SID_SYSNR
```

3. Change the file `j2ee-install-dir/server/services/security/work/R3Security.properties` as follows:

```
sapbasis.ashost = logical-host-as
```

▼ *How to Configure a SAP Web Dispatcher with your Sun Cluster HA for SAP*

After you have configured the SAP Web dispatcher with your Sun Cluster HA for SAP, perform the following steps.

1. **Ensure that SAP Web dispatcher has an instance number different than the Central Instance and the application server instances.**
For example, SAPSYSTEM = 66 is used in the profile for the SAP Web dispatcher.
2. **Activate the Internet Communication Frame Services manually after you install the SAP Web Application Server.**
See SAP OSS note 517484 for more details.

Configuring Sun ONE Web Server

The “How to Configure a Sun ONE Web Server” procedure in the Sun Cluster data services collection is missing the following step, which is not dependent on any other step in the procedure.

Create a file that contains the secure key password you need to start this instance, and place this file under the server root directory. Name this file `keypass`.

Note – Because this file contains the key database password, protect the file with the appropriate permissions.

Support for Sun ONE Proxy Server

Sun Cluster HA for Sun ONE Web Server now supports Sun ONE Proxy Server. For information about the Sun ONE Proxy Server product, see <http://docs.sun.com/db/prod/s1.webproxys>. For Sun ONE Proxy Server installation and configuration information, see http://docs.sun.com/db/coll/S1_ipwebproxysrvr36.

Registering and Configuring the Sun Cluster for SAP liveCache

The procedure “How to Register and Configure Sun Cluster HA for SAP liveCache” has been revised. Add the following command to step 4 of this procedure.

```
-x affinityon=TRUE
```

Note – `AffinityOn` must be set to `TRUE` and the local file system must reside on global disk groups to be failover

For the procedure on how to set up an `HAStoragePlus` resource, see *Sun Cluster 3.0 Data Service Installation and Configuration Guide*.

Using the Sun Cluster `LogicalHostname` Resource With Oracle Parallel Server/Real Application Clusters

Information on using the Sun Cluster `LogicalHostname` resource with Oracle Parallel Server/Real Application Clusters is missing from “Installing and Configuring Sun Cluster Support for Oracle Parallel Server/Real Application Clusters” in *Sun Cluster 3.0 12/01 Data Services Installation and Configuration Guide*.

If a cluster node that is running an instance of Oracle Parallel Server/Real Application Clusters fails, an operation that a client application attempted might be required to time out before the operation is attempted again on another instance. If the TCP/IP network timeout is high, the client application might take a long time to detect the failure. Typically client applications take between three and nine minutes to detect such failures.

In such situations, client applications may use the Sun Cluster `LogicalHostname` resource for connecting to an Oracle Parallel Server/Real Application Clusters database that is running on Sun Cluster. You can configure the `LogicalHostname` resource in a separate resource group that is mastered on the nodes on which Oracle Parallel Server/Real Application Clusters is running. If a node fails, the `LogicalHostname` resource fails over to another surviving node on which Oracle Parallel Server/Real Application Clusters is running. The failover of the `LogicalHostname` resource enables new connections to be directed to the other instance of Oracle Parallel Server/Real Application Clusters.



Caution – Before using the `LogicalHostname` resource for this purpose, consider the effect on existing user connections of failover or failback of the `LogicalHostname` resource.

Creating Node-Specific Files and Directories for Use With Oracle Parallel Server/Real Application Clusters Software on the Cluster File System

When Oracle software is installed on the cluster file system, all the files in the directory that the `ORACLE_HOME` environment variable specifies are accessible by all cluster nodes.

An installation might require that some Oracle files or directories maintain node-specific information. You can satisfy this requirement by using a symbolic link whose target is a file or a directory on a file system that is local to a node. Such a file system is not part of the cluster file system.

To use a symbolic link for this purpose, you must allocate an area on a local file system. To enable Oracle applications to create symbolic links to files in this area, the applications must be able to access files in this area. Because the symbolic links reside on the cluster file system, all references to the links from all nodes are the same. Therefore, all nodes must have the same namespace for the area on the local file system.

▼ *How to Create a Node-Specific Directory for Use With Oracle Parallel Server/Real Application Clusters Software on the Cluster File System*

Perform this procedure for each directory that is to maintain node-specific information. The following directories are typically required to maintain node-specific information:

- \$ORACLE_HOME/network/agent
- \$ORACLE_HOME/network/log
- \$ORACLE_HOME/network/trace
- \$ORACLE_HOME/srvm/log
- \$ORACLE_HOME/apache

For information about other directories that might be required to maintain node-specific information, see your Oracle documentation.

1. On each cluster node, create the local directory that is to maintain node-specific information.

```
# mkdir -p local-dir
```

-p Specifies that all nonexistent parent directories are created first

local-dir Specifies the full path name of the directory that you are creating

2. On each cluster node, make a local copy of the global directory that is to maintain node-specific information.

```
# cp -pr global-dir local-dir-parent
```

-p Specifies that the owner, group, permissions modes, modification time, access time, and access control lists are preserved.

-r Specifies that the directory and all its files, including any subdirectories and their files, are copied.

global-dir Specifies the full path of the global directory that you are copying. This directory resides on the cluster file system under the directory that the ORACLE_HOME environment variable specifies.

local-dir-parent Specifies the directory on the local node that is to contain the local copy. This directory is the parent directory of the directory that you created in Step 1.

3. Replace the global directory that you copied in Step 2 with a symbolic link to the local copy of the global directory.

a. From any cluster node, remove the global directory that you copied in Step 2.

```
# rm -r global-dir
```


-r Specifies that the directory and all its files, including any subdirectories and their files, are removed.

global-dir Specifies the file name and full path of the global directory that you are removing. This directory is the global directory that you copied in Step 2.

b. From any cluster node, create a symbolic link from the local copy of the directory to the global directory that you removed in Step a.

```
# ln -s local-dir global-dir
```

-s Specifies that the link is a symbolic link

local-dir Specifies that the local directory that you created in Step 1 is the source of the link

global-dir Specifies that the global directory that you removed in Step a is the target of the link

EXAMPLE 3-1 Creating Node-Specific Directories

This example shows the sequence of operations that is required to create node-specific directories on a two-node cluster. This cluster is configured as follows:

- The ORACLE_HOME environment variable specifies the /global/oracle directory.
- The local file system on each node is located under the /local directory.

The following operations are performed on each node:

1. To create the required directories on the local file system, the following commands are run:

```
# mkdir -p /local/oracle/network/agent
# mkdir -p /local/oracle/network/log
# mkdir -p /local/oracle/network/trace
# mkdir -p /local/oracle/srvm/log
# mkdir -p /local/oracle/apache
```

2. To make local copies of the global directories that are to maintain node-specific information, the following commands are run:

```
# cp -pr $ORACLE_HOME/network/agent /local/oracle/network/.
# cp -pr $ORACLE_HOME/network/log /local/oracle/network/.
# cp -pr $ORACLE_HOME/network/trace /local/oracle/network/.
# cp -pr $ORACLE_HOME/srvm/log /local/oracle/srvm/.
# cp -pr $ORACLE_HOME/apache /local/oracle/.
```

The following operations are performed on only one node:

EXAMPLE 3-1 Creating Node-Specific Directories (Continued)

1. To remove the global directories, the following commands are run:

```
# rm -r $ORACLE_HOME/network/agent
# rm -r $ORACLE_HOME/network/log
# rm -r $ORACLE_HOME/network/trace
# rm -r $ORACLE_HOME/srvm/log
# rm -r $ORACLE_HOME/apache
```

2. To create symbolic links from the local directories to their corresponding global directories, the following commands are run:

```
# ln -s /local/oracle/network/agent $ORACLE_HOME/network/agent
# ln -s /local/oracle/network/log $ORACLE_HOME/network/log
# ln -s /local/oracle/network/trace $ORACLE_HOME/network/trace
# ln -s /local/oracle/srvm/log $ORACLE_HOME/srvm/log
# ln -s /local/oracle/apache $ORACLE_HOME/apache
```

▼ *How to Create a Node-Specific File for Use With Oracle Parallel Server/Real Application Clusters Software on the Cluster File System*

Perform this procedure for each file that is to maintain node-specific information. The following files are typically required to maintain node-specific information:

- \$ORACLE_HOME/network/admin/snmp_ro.ora
- \$ORACLE_HOME/network/admin/snmp_rw.ora

For information about other files that might be required to maintain node-specific information, see your Oracle documentation.

1. **On each cluster node, create the local directory that will contain the file that is to maintain node-specific information.**

```
# mkdir -p local-dir
```

-p Specifies that all nonexistent parent directories are created first

local-dir Specifies the full path name of the directory that you are creating

2. **On each cluster node, make a local copy of the global file that is to maintain node-specific information.**

```
# cp -p global-file local-dir
```

-p Specifies that the owner, group, permissions modes, modification time, access time, and access control lists are preserved.

- global-file* Specifies the file name and full path of the global file that you are copying. This file was installed on the cluster file system under the directory that the ORACLE_HOME environment variable specifies.
- local-dir* Specifies the directory that is to contain the local copy of the file. This directory is the directory that you created in Step 1.

3. Replace the global file that you copied in Step 2 with a symbolic link to the local copy of the file.

a. From any cluster node, remove the global file that you copied in Step 2.

```
# rm global-file
```

global-file Specifies the file name and full path of the global file that you are removing. This file is the global file that you copied in Step 2.

b. From any cluster node, create a symbolic link from the local copy of the file to the directory from which you removed the global file in Step a.

```
# ln -s local-file global-dir
```

-s Specifies that the link is a symbolic link

local-file Specifies that the file that you copied in Step 2 is the source of the link

global-dir Specifies that the directory from which you removed the global version of the file in Step a is the target of the link

EXAMPLE 3-2 Creating Node-Specific Files

This example shows the sequence of operations that is required to create node-specific files on a two-node cluster. This cluster is configured as follows:

- The ORACLE_HOME environment variable specifies the /global/oracle directory.
- The local file system on each node is located under the /local directory.

The following operations are performed on each node:

1. To create the local directory that will contain the files that are to maintain node-specific information, the following command is run:

```
# mkdir -p /local/oracle/network/admin
```

2. To make a local copy of the global files that are to maintain node-specific information, the following commands are run:

```
# cp -p $ORACLE_HOME/network/admin/snmp_ro.ora \
  /local/oracle/network/admin/.
```

```
# cp -p $ORACLE_HOME/network/admin/snmp_rw.ora \
  /local/oracle/network/admin/.
```

The following operations are performed on only one node:

EXAMPLE 3-2 Creating Node-Specific Files (Continued)

1. To remove the global files, the following commands are run:

```
# rm $ORACLE_HOME/network/admin/snmp_ro.ora
# rm $ORACLE_HOME/network/admin/snmp_rw.ora
```

2. To create symbolic links from the local copies of the files to their corresponding global files, the following commands are run:

```
# ln -s /local/oracle/network/admin/snmp_ro.ora \
$ORACLE_HOME/network/admin/snmp_rw.ora
# ln -s /local/oracle/network/admin/snmp_rw.ora \
$ORACLE_HOME/network/admin/snmp_rw.ora
```

Supplement

The following subsections describe known errors in or omissions from the *Sun Cluster 3.0 5/02 Supplement*.

How to Uninstall Sun Cluster Software From a Cluster Node (5/02)

The following note at the beginning of this procedure is incorrect:

Note – To uninstall Sun Cluster software from a node that has not yet joined the cluster or is still in install mode, do not perform this procedure. Instead, go to "How to Uninstall Sun Cluster Software to Correct Installation Problems" in the *Sun Cluster 3.0 12/01 Software Installation Guide*.

The note should instead read as follows:

Note – To uninstall Sun Cluster software from a node that has not yet joined the cluster or is still in install mode, do not perform this procedure. Instead, go to "How to Uninstall Sun Cluster Software to Correct Installation Problems" in the *Sun Cluster 3.0 5/02 Supplement*.

Release Notes

The following subsections describe omissions or new information that will be added to the next publishing of the Release Notes.

BugId 4662264

The Workaround documented in the *Sun Cluster 3.0 5/02 Release Notes* is incorrect.

Incorrect:

Increase the stack size by putting the following lines in the `/etc/system` file.

```
set lwp_default_stksize=0x6000
set svc_default_stksize 0x8000
```

Correct:

Increase the stack size by putting the following lines in the `/etc/system` file.

```
set lwp_default_stksize=0x6000
set rpcmod:svc_default_stksize=0x8000
```

Man Pages

The following subsections describe omissions or new information that will be added to the next publishing of the man pages.

scconf_transp_adap_sci Man Page

The `scconf_transp_adap_sci(1M)` man page states that SCI transport adapters can be used with the `rsm` transport type. This support statement is incorrect. SCI transport adapters do *not* support the `rsm` transport type. SCI transport adapters support the `dlpi` transport type only.

scconf_transp_adap_wrsm Man Page

The following `scconf_transp_adap_wrsm(1M)` man page replaces the existing `scconf_transp_adap_wrsm(1M)` man page.

NAME

`scconf_transp_adap_wrsm.1m`- configure the `wrsm` transport adapter

DESCRIPTION

`wrsm` adapters may be configured as cluster transport adapters. These adapters can only be used with transport types `dlpi`.

The `wrsm` adapter connects to a transport junction or to another `wrsm` adapter on a different node. In either case, the connection is made through a transport cable.

Although you can connect the `wrsm` adapters directly by using a point-to-point configuration, Sun Cluster software requires that you specify a transport junction, a virtual transport junction. For example, if `node1:wrsm1` is connected to `node2:wrsm1` directly through a cable, you must specify the following configuration information.

```
node1:wrsm1 <--cable1--> Transport Junction sw_wrsm1 <--cable2--> node2:wrsm1
```

The transport junction, whether a virtual switch or a hardware switch, must have a specific name. The name must be `sw_wrsmN` where the adapter is `wrsmN`. This requirement reflects a Wildcat restriction that requires that all `wrsm` controllers on the same Wildcat network have the same instance number.

When a transport junction is used and the endpoints of the transport cable are configured using `sconf`, `scinstall`, or other tools, you are asked to specify a port name on the transport junction. You can provide any port name, or accept the default, as long as the name is unique for the transport junction.

The default sets the port name to the node ID that hosts the adapter at the other end of the cable.

Refer to `sconf(1M)` for more configuration details.

There are no user configurable properties for cluster transport adapters of this type.

SEE ALSO

`sconf(1M)`, `scinstall(1M)`, `wrsmconf(1M)`, `wrsmstat(1M)`, `wrsm(7D)`, `wrsm(7D)`

SUNW.HAStoragePlus.5

The `SUNW.HAStoragePlus.5` man page has been modified. The following paragraph replaces the paragraph in the Notes section of the man page.

Although unlikely, the `SUNW.HAStoragePlus` resource is capable of mounting any global file system found to be in a unmounted state. This check will be skipped only if the file system is of type `UFS` and logging is turned off. All file systems are mounted in the overlay mode. Local file systems will be forcibly unmounted.

The following `FilesystemCheckCommand` extension property has been added to the `SUNW.HAStoragePlus.5` man page.

FilesystemCheckCommand

SUNW.HAStoragePlus conducts a file system check on each unmounted file system before attempting to mount it. The default file system check command is `/usr/sbin/fsck -o p` for UFS and VxFS file systems, and `/usr/sbin/fsck` for other file systems. The `FilesystemCheckCommand` extension property can be used to override this default file system check specification and instead specify an alternate command string/executable. This command string/executable will then be invoked on all unmounted file systems.

The default `FilesystemCheckCommand` extension property value is `NULL`. When the `FilesystemCheckCommand` is set to `NULL` the command will be assumed to be `/usr/sbin/fsck -o p` for UFS/VxFS file systems and `/usr/sbin/fsck` for other file systems. When the `FilesystemCheckCommand` is set to a user specified command string, `SUNW.HAStoragePlus` will elect to invoke this command string with the file system mount point as an argument. Any arbitrary executable can be specified in this manner. A non-zero return value will be treated as a error which occurred during the file system check operation, causing the start method to fail. Any arbitrary executable can be specified in this manner. When the `FilesystemCheckCommand` is set to `/usr/bin/true`, file system checks will altogether be avoided.

Scalable Cluster Topology

This appendix provides information and procedures for using the scalable cluster topology. This information supplements the . Certain procedures have been updated and included here to accommodate this new Sun Cluster 3.x topology.

This chapter contains new information for the following topics.

- “Overview of Scalable Topology” on page 57
- “Adding a Cluster Node” on page 58
- “Removing a Cluster Node” on page 58
- “How to Remove Connectivity Between an Array and a Single Node, in a Cluster With Greater Than Two-Node Connectivity” on page 60

Overview of Scalable Topology

The scalable cluster topology allows connectivity of up to four nodes to a single storage array. Note the following considerations for this topology at this time:

- All nodes must have the Oracle Parallel Server/Real Application Clusters (OPS/RAC) software installed. For information about installing and using OPS/RAC in a cluster, see the *Sun Cluster 3.0 12/01 Data Services Installation and Configuration Guide*.
- The storage arrays supported with this cluster topology include the Sun StorEdge T3/T3+ array (single-controller and partner-group configurations), the Sun StorEdge 9900 Series storage device, and the Sun StorEdge 3900 storage device.

Adding or Removing a Cluster Node

The following information and procedures supplement procedures in the .

Adding a Cluster Node

The scalable topology does not introduce any changes to the standard procedure for adding cluster nodes. See Sun Cluster system administration documentation for the procedure for adding a cluster node.

Figure A-1 shows a sample diagram of cabling for four-node connectivity with scalable topology.

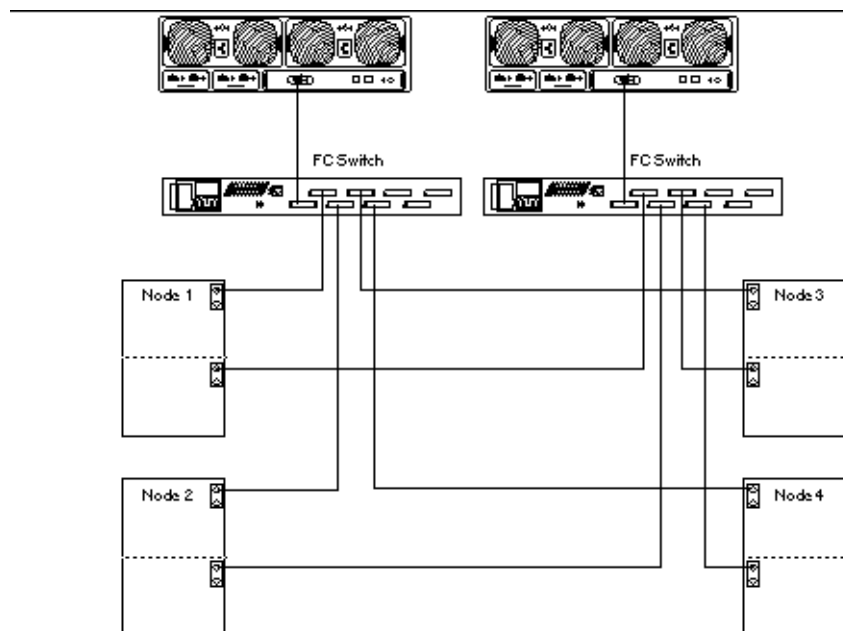


FIGURE A-1 Sample Scalable Topology Cabling, Four-Node Connectivity

Removing a Cluster Node

The following task map procedure is an update to the standard procedure in the Sun Cluster system administration documentation.



Caution – Do not use this procedure if your cluster is running an OPS configuration. At this time, removing a node in an OPS configuration might cause nodes to panic at reboot.

TABLE A-1 Task Map: Removing a Cluster Node (5/02)

Task	For Instructions, Go To
Move all resource groups and disk device groups off of the node to be removed. - Use <code>scswitch</code>	<code># scswitch -S -h from-node</code>
Remove the node from all resource groups. - Use <code>scrgadm</code>	Sun Cluster data services collection: See the procedure for how to remove a node from an existing resource group.
Remove node from all disk device groups - Use <code>scconf</code> , <code>metaset</code> , and <code>scsetup</code>	Sun Cluster system administration documentation: see the procedures for how to remove a node from a disk device group (separate procedures for Solstice DiskSuite, VERITAS Volume Manager, and raw disk device groups).
Remove all quorum devices. - Use <code>scsetup</code> .	<p>Caution: Do not remove the quorum device if you are removing a node from a two-node cluster.</p> Sun Cluster system administration documentation: "How to Remove a Quorum Device." Note that although you must remove the quorum device before you remove the storage device in the next step, you can add the quorum device back immediately afterward.
Remove the storage device from the node. - Use <code>devfsadm</code> , <code>scdidadm</code> .	"How to Remove Connectivity Between an Array and a Single Node, in a Cluster With Greater Than Two-Node Connectivity" on page 60
Add the new quorum device (to only the nodes that are intended to remain in the cluster). - Use <code>scconf -a -q globaldev=d[n], node=node1, node=node2, . . .</code>	<code>scconf(1M)</code> man page
Place the node being removed into maintenance state. - Use <code>scswitch</code> , <code>shutdown</code> , and <code>scconf</code> .	Sun Cluster system administration documentation: "How to Put a Node Into Maintenance State"

TABLE A-1 Task Map: Removing a Cluster Node (5/02) (Continued)

Task	For Instructions, Go To
Remove all logical transport connections to the node being removed. - Use <code>scsetup</code> .	Sun Cluster system administration documentation: "How to Remove Cluster Transport Cables, Transport Adapters, and Transport Junctions"
Remove node from the cluster software configuration. - Use <code>scconf</code> .	Sun Cluster system administration documentation: "How to Remove a Node From the Cluster Software Configuration"

▼ How to Remove Connectivity Between an Array and a Single Node, in a Cluster With Greater Than Two-Node Connectivity

Use this procedure to detach a storage array from a single cluster node, in a cluster that has three- or four-node connectivity.

1. **Back up all database tables, data services, and volumes that are associated with the storage array that you are removing.**
2. **Determine the resource groups and device groups that are running on the node to be disconnected.**

`scstat`
3. **If necessary, move all resource groups and device groups off the node to be disconnected.**



Caution – If your cluster is running OPS/RAC software, shut down the OPS/RAC database instance that is running on the node before you move the groups off the node. For instructions see the *Oracle Database Administration Guide*.

```
# scswitch -S -h from-node
```

4. **Put the device groups into maintenance state.**
For the procedure on quiescing I/O activity to Veritas shared disk groups, see your VERITAS Volume Manager documentation.
For the procedure on putting a device group in maintenance state, see the Sun Cluster system administration documentation.
5. **Remove the node from the device groups.**

- If you use VERITAS Volume Manager or raw disk, use the `scconf` command to remove the device groups.
 - If you use Solstice DiskSuite/Solaris Volume Manager, use the `metaset` command to remove the device groups.
6. **If the cluster is running HAStorage or HAStoragePlus, remove the node from the resource group's nodelist.**
- ```
scrgadm -a -g resource-group -h nodelist
```
- See the Sun Cluster data services collection for more information on changing a resource group's nodelist.
7. **If the storage array you are removing is the last storage array that is connected to the node, disconnect the fiber-optic cable between the node and the hub or switch that is connected to this storage array (otherwise, skip this step).**
8. **Do you want to remove the host adapter from the node you are disconnecting?**
- If yes, shut down and power off the node.
  - If no, skip to Step 11.
9. **Remove the host adapter from the node.**
- For the procedure on removing host adapters, see the documentation that shipped with your node.
10. **Without allowing the node to boot, power on the node.**
- For more information, see the Sun Cluster system administration documentation.
11. **Boot the node into non-cluster mode.**

```
ok boot -x
```




---

**Caution** – The node must be in non-cluster mode before you remove OPS/RAC software in the next step or the node will panic and potentially cause a loss of data availability.

---

12. **If OPS/RAC software has been installed, remove the OPS/RAC software package from the node that you are disconnecting.**

```
pkgrm SUNWscum
```




---

**Caution** – If you do not remove the OPS/RAC software from the node you disconnected, the node will panic when the node is reintroduced to the cluster and potentially cause a loss of data availability.

---

13. **Boot the node into cluster mode.**

```
ok> boot
```

For more information, see the Sun Cluster system administration documentation.

**14. On the node, update the device namespace by updating the /devices and /dev entries.**

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

**15. Bring the device groups back online.**

For procedures on bringing a VERITAS shared disk group online, see your VERITAS Volume Manager documentation.

For the procedure on bringing a device group online, see the procedure on putting a device group into maintenance state in the Sun Cluster system administration documentation.

# Installing and Configuring Sun Cluster HA for SAP liveCache

---

This chapter contains the procedures on how to install and configure Sun Cluster HA for SAP liveCache.

This chapter contains the following procedures.

- “How to Prepare the Nodes” on page 69
- “How to Install and Configure liveCache” on page 70
- “How to Enable liveCache to Run in a Cluster” on page 70
- “How to Verify the liveCache Installation and Configuration” on page 71
- “How to Install the Sun Cluster HA for SAP liveCache Packages” on page 72
- “How to Register and Configure Sun Cluster HA for SAP liveCache” on page 75
- “How to Verify the Sun Cluster HA for SAP liveCache Installation and Configuration” on page 78

---

## Sun Cluster HA for SAP liveCache Overview

Use the information in this section to understand how Sun Cluster HA for SAP liveCache makes liveCache highly available.


For conceptual information on failover and scalable services, see the *Sun Cluster 3.1 Concepts Guide*.

To eliminate a single point of failure in an SAP Advanced Planner & Optimizer (APO) System, Sun Cluster HA for SAP liveCache provides fault monitoring and automatic failover for liveCache and fault monitoring and automatic restart for SAP xserver. The following table lists the data services that best protect SAP Supply Chain Management (SCM) components in a Sun Cluster configuration. Figure B-1 also illustrates the data services that best protect SAP SCM components in a Sun Cluster configuration.

**TABLE B-1** Protection of liveCache Components

| <b>liveCache Component</b> | <b>Protected by</b>                                                                                                                                                  |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SAP APO Central Instance   | Sun Cluster HA for SAP<br>The resource type is <code>SUNW.sap_ci_v2</code> .<br>For more information on this data service, see Sun Cluster data services collection. |
| SAP APO database           | All highly available databases that are supported with Sun Cluster software and by SAP.                                                                              |
| SAP APO Application Server | Sun Cluster HA for SAP<br>The resource type is <code>SUNW.sap_as_v2</code> .<br>For more information on this data service, see Sun Cluster data services collection. |
| SAP liveCache xserver      | Sun Cluster HA for SAP liveCache<br>The resource type is <code>SUNW.sap_xserver</code> .                                                                             |
| SAP liveCache database     | Sun Cluster HA for SAP liveCache<br>The resource type is <code>SUNW.sap_livecache</code> .                                                                           |
| NFS file system            | Sun Cluster HA for NFS<br>The resource type is <code>SUNW.nfs</code> .<br>For more information on this data service, see Sun Cluster data services collection.       |





**FIGURE B-1** Protection of liveCache Components

---

## Installing and Configuring Sun Cluster HA for SAP liveCache

Table B-2 lists the tasks for installing and configuring Sun Cluster HA for SAP liveCache. Perform these tasks in the order that they are listed.

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

| Task                                                   | For Instructions, Go To                                         |
|--------------------------------------------------------|-----------------------------------------------------------------|
| Plan the Sun Cluster HA for SAP liveCache installation | Your SAP documentation.<br>Sun Cluster data services collection |
| Prepare the nodes and disks                            | "How to Prepare the Nodes" on page 69                           |

**TABLE B-2** Task Map: Installing and Configuring Sun Cluster HA for SAP liveCache  
(Continued)

| Task                                                                               | For Instructions, Go To                                                                                         |
|------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| Install and configure liveCache                                                    | "How to Install and Configure liveCache" on page 70<br>"How to Enable liveCache to Run in a Cluster" on page 70 |
| Verify liveCache installation and configuration                                    | "How to Verify the liveCache Installation and Configuration" on page 71                                         |
| Install Sun Cluster HA for SAP liveCache packages                                  | "How to Install the Sun Cluster HA for SAP liveCache Packages" on page 72                                       |
| Register and configure Sun Cluster HA for SAP liveCache as a failover data service | "How to Register and Configure Sun Cluster HA for SAP liveCache" on page 75                                     |
| Verify Sun Cluster HA for SAP liveCache installation and configuration             | "Verifying the Sun Cluster HA for SAP liveCache Installation and Configuration" on page 78                      |
| Understand Sun Cluster HA for SAP liveCache Fault Monitors                         | "Understanding Sun Cluster HA for SAP liveCache Fault Monitors" on page 80                                      |

---

## Planning the Sun Cluster HA for SAP liveCache Installation and Configuration

This section contains the information you need to plan your Sun Cluster HA for SAP liveCache installation and configuration.

---

**Note** – If you have not already done so, read your SAP documentation before you begin planning your Sun Cluster HA for SAP liveCache installation and configuration because your SAP documentation includes configuration restrictions and requirements that are not outlined in Sun Cluster documentation or dictated by Sun Cluster software.

---

## Configuration Requirements



---

**Caution** – Your data service configuration might not be supported if you do not adhere to these requirements.

---

Use the requirements in this section to plan the installation and configuration of Sun Cluster HA for SAP liveCache. These requirements apply to Sun Cluster HA for SAP liveCache only. You must meet these requirements before you proceed with your Sun Cluster HA for SAP liveCache installation and configuration.


For requirements that apply to all data services, see Sun Cluster data services collection.

- Use an SAP liveCache version 7.4 or higher.
- Configure SAP xserver so that SAP xserver starts on all nodes that the liveCache resource can failover to. To implement this configuration, ensure that the `nodelist` of the SAP xserver resource group and the liveCache resource group contain the same nodes. Also, the value of `desired primaries` and `maximum primaries` of the SAP xserver resource must be equal to the number of nodes listed in the `nodelist` parameter of the liveCache resource. For more information, see “How to Register and Configure Sun Cluster HA for SAP liveCache” on page 75.

## Standard Data Service Configurations

Use the standard configurations in this section to plan the installation and configuration of Sun Cluster HA for SAP liveCache. Sun Cluster HA for SAP liveCache supports the standard configurations in this section. Sun Cluster HA for SAP liveCache might support additional configurations. However, you must contact your Sun service provider for information on additional configurations.

Figure B-2 illustrates a four-node cluster with SAP APO Central Instance, APO application servers, a database, and liveCache. APO Central Instance, the database, and liveCache are configured as failover data services. APO application servers and SAP xserver can be configured as scalable or failover data services.



**FIGURE B-2** Four-Node Cluster

## Configuration Considerations

Use the information in this section to plan the installation and configuration of Sun Cluster HA for SAP liveCache. The information in this section encourages you to think about the impact your decisions have on the installation and configuration of Sun Cluster HA for SAP liveCache.

- Install liveCache on its own global device group, separate from the global device group for the APO Oracle database and SAP R/3 software. This separate global device group for liveCache ensures that the liveCache resource can depend on the `HASStoragePlus` resource for liveCache only.
- If you want to run SAP xserver as any user other than user `root`, create that user on all nodes on which SAP xserver runs, and define this user in the `Xserver_User` extension property. SAP xserver starts and stops based on the user you identify in this extension property. The default for this extension property is user `root`.
- Configure SAP xserver as a failover resource unless you are running multiple liveCache instances that overlap.

## Configuration Planning Questions

Use the questions in this section to plan the installation and configuration of Sun Cluster HA for SAP liveCache. Insert the answers to these questions into the data service worksheets in the *Sun Cluster 3.1 Release Notes*. See "Configuration Considerations" on page 68 for information that might apply to these questions.

- What resource groups will you use for network addresses and application resources and the dependencies between them?
- What is the logical hostname (for liveCache resource) for clients that will access the data service?
- Where will the system configuration files reside?  
See Sun Cluster data services collection for the advantages and disadvantages of placing the liveCache binaries on the local file system as opposed to the cluster file system.

---

## Preparing the Nodes and Disks

This section contains the procedures you need to prepare the nodes and disks.

### ▼ How to Prepare the Nodes

Use this procedure to prepare for the installation and configuration of liveCache.

1. **Become superuser on all of the nodes.**
2. **Configure the `/etc/nsswitch.conf` file.**
  - a. **On each node that can master the liveCache resource, include one of the following entries for `group`, `project`, an `passwd` database entries in the `/etc/nsswitch.conf` file.**

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

- b. **On each node that can master the liveCache resource, ensure that `files` appears first for the `protocols` database entry in the `/etc/nsswitch.conf` file.**

Example:

```
protocols: files nis
```

Sun Cluster HA for SAP liveCache uses the `su - user` command and the `dbmcli` command to start and stop liveCache.

The network information name service might become unavailable when a cluster node's public network fails. Implementing the preceding changes to the `/etc/nsswitch.conf` file ensures that the `su(1M)` command and the `dbmcli` command do not refer to the NIS/NIS+ name services.

---

## Installing and Configuring liveCache

This section contains the procedures you need to install and configure liveCache.

## ▼ How to Install and Configure liveCache

Use this procedure to install and configure liveCache.

### 1. Install and configure SAP APO System.

See Sun Cluster data services collection for the procedures on how to install and configure SAP APO System on Sun Cluster software.

### 2. Install liveCache.

---

**Note** – Install liveCache by using the physical hostname if you have not already created the required logical host.

---

For more information, see your SAP documentation.

### 3. Create the `.XUSER.62` file for the SAP APO administrator user and the liveCache administrator user by using the following command.

```
dbmcli -d LC-NAME -n logical-hostname -us user,passwd
```

|                  |                                                           |
|------------------|-----------------------------------------------------------|
| LC-NAME          | Uppercase name of liveCache database instance             |
| logical-hostname | Logical hostname that is used with the liveCache resource |



---

**Caution** – Neither SAP APO transaction LC10 nor Sun Cluster HA for SAP liveCache functions properly if you do not create this file correctly.

---

### 4. Copy `/usr/spool/sql` from the node, on which you installed liveCache, to all the nodes that will run the liveCache resource. Ensure that the ownership of these files is the same on all node as it is on the node on which you installed liveCache.

Example:

```
tar cfb - /usr/spool/sql | rsh phys-schost-1 tar xfb -
```

## ▼ How to Enable liveCache to Run in a Cluster

During a standard SAP installation, liveCache is installed with a physical hostname. You must modify liveCache to use a logical hostname so that liveCache works in a Sun Cluster environment. Use this procedure to enable liveCache to run in a cluster.

### 1. Create the failover resource group to hold the network and liveCache resource.

```
scrgadm -a -g livecache-resource-group [-h nodelist]
```

2. Verify that you added all the network resources you use to your name service database.

3. Add a network resource (logical hostname) to the failover resource group.

```
scrgadm -a -L -g livecache-resource-group -l lc-logical-hostname [-n netiflist]
```

4. Enable the failover resource group.

```
scswitch -Z -g livecache-resource-group
```

5. Log on to the node that hosts the liveCache resource group.

6. Start SAP xserver manually on the node that hosts the liveCache resource group.

```
su - lc-nameadm
x_server start
```

lc-name Lowercase name of liveCache database instance

7. Log on to SAP APO System by using your SAP GUI with user DDIC.

8. Go to transaction LC10 and change the liveCache host to the logical hostname you defined in Step 3.

```
liveCache host: lc-logical-hostname
```

---

## Verifying the liveCache Installation and Configuration

This section contains the procedure you need to verify the liveCache installation and configuration.

### ▼ How to Verify the liveCache Installation and Configuration

Use this procedure to verify the liveCache installation and configuration. This procedure does not verify that your application is highly available because you have not installed your data service yet.

1. Log on to SAP APO System by using your SAP GUI with user DDIC.

2. Go to transaction LC10.
3. Ensure that you can check the state of liveCache.
4. Ensure that the following `dbmcli` command work as user `lc_nameadm`.

```
dbmcli -d LC_NAME -n logical-hostname db_state
dbmcli -d LC_NAME -n logical-hostname db_enum
```

---

## Installing the Sun Cluster HA for SAP liveCache Packages

This section contains the procedure you need to install the Sun Cluster HA for SAP liveCache packages.

### ▼ How to Install the Sun Cluster HA for SAP liveCache Packages

Use this procedure to install the Sun Cluster HA for SAP liveCache packages. You need the Sun Cluster 3.1 Agents CD-ROM to perform this procedure. This procedure assumes that you did not install the data service packages during your initial Sun Cluster installation.

1. **Load the Sun Cluster 3.1 Agents CD-ROM into the CD-ROM drive.**
2. **Run the `scinstall` utility with no options.**  
This step starts the `scinstall` utility in interactive mode.
3. **Choose the Add Support for New Data Service to This Cluster Node menu option.**  
The `scinstall` utility prompts you for additional information.
4. **Provide the path to the Sun Cluster 3.1 Agents CD-ROM.**  
The utility refers to the CD-ROM as the "data services cd."
5. **Specify the data service to install.**  
The `scinstall` utility lists the data service that you selected and asks you to confirm your choice.
6. **Exit the `scinstall` utility.**



7. Unload the CD-ROM from the drive.

---

## Registering and Configuring the Sun Cluster HA for SAP liveCache

This section contains the procedures you need to configure Sun Cluster HA for SAP liveCache.

### Sun Cluster HA for SAP liveCache Extension Properties

Use the extension properties in Table B-3 and Table B-4 to create your resources. Use the following command line to configure extension properties when you create your resource.

```
scrgadm -x parameter=value
```

Use the procedure in Sun Cluster data services collection to configure the extension properties if you have already created your resources. You can update some extension properties dynamically. You can update others, however, only when you create or disable a resource. The Tunable fields in Table B-3 and Table B-4 indicate when you can update each property. See Appendix A for details on all Sun Cluster properties.

**TABLE B-3** Sun Cluster HA for SAP liveCache (SUNW.sap\_xserver) Extension Properties

| Name/Data Type                 | Description                                                                                                                                   |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| Confdir_List (optional) String | The directory for liveCache software and instance directories.<br><b>Default:</b> /sapdb<br><b>Range:</b> None<br><b>Tunable:</b> At creation |
| Monitor_retry_count            | Number of PMF restarts that are allowed for the fault monitor.<br><b>Default:</b> 4<br><b>Tunable:</b> Any time                               |

**TABLE B-3** Sun Cluster HA for SAP liveCache (SUNW.sap\_xserver) Extension Properties (Continued)

| Name/Data Type                   | Description                                                                                                                                                                                                                                |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Monitor_retry_interval           | Time interval in minutes for fault monitor restarts.<br><b>Default:</b> 2<br><b>Tunable:</b> Any time                                                                                                                                      |
| Probe_timeout                    | Time-out value in seconds for the probes.<br><b>Default:</b> 120<br><b>Tunable:</b> Any time                                                                                                                                               |
| Soft_Stop_Pct (optional) Integer | Percentage of stop timeout that is used to stop SAP xserver by using the SAP utility x_server stop before SIGKILL is used to stop all SAP xserver processes.<br><b>Default:</b> 50<br><b>Range:</b> 1-100<br><b>Tunable:</b> When disabled |
| Xserver_User (optional) String   | SAP xserver system administrator user name.<br><b>Default:</b> root<br><b>Range:</b> None<br><b>Tunable:</b> At creation                                                                                                                   |

**TABLE B-4** Sun Cluster HA for SAP liveCache (SUNW.sap\_livecache) Extension Properties

| Name/Data Type                   | Description                                                                                                                                     |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| Confdir_list (optional) String   | The directory for liveCache software and the instance directory.<br><b>Default:</b> /sapdb<br><b>Range:</b> None<br><b>Tunable:</b> At creation |
| Livecache_name (required) String | Name of liveCache database instance.<br><b>Default:</b> None<br><b>Range:</b> None<br><b>Tunable:</b> At creation                               |

**TABLE B-4** Sun Cluster HA for SAP liveCache (SUNW.sap\_livecache) Extension Properties (Continued)

| Name/Data Type         | Description                                                                                                     |
|------------------------|-----------------------------------------------------------------------------------------------------------------|
| Monitor_retry_count    | Number of PMF restarts that are allowed for the fault monitor.<br><b>Default:</b> 4<br><b>Tunable:</b> Any time |
| Monitor_retry_interval | Time interval in minutes for fault monitor restarts.<br><b>Default:</b> 2<br><b>Tunable:</b> Any time           |
| Probe_timeout          | Time-out value in seconds for the probes.<br><b>Default:</b> 90<br><b>Tunable:</b> Any time                     |

## ▼ How to Register and Configure Sun Cluster HA for SAP liveCache

Use this procedure to configure Sun Cluster HA for SAP liveCache as a failover data service for the liveCache database and SAP xserver as a failover or scalable data service. This procedure assumes that you installed the data service packages. If you did not install the Sun Cluster HA for SAP liveCache packages as part of your initial Sun Cluster installation, go to “How to Install the Sun Cluster HA for SAP liveCache Packages” on page 72 to install the data service packages. Otherwise, use this procedure to configure the Sun Cluster HA for SAP liveCache.




---

**Caution** – Do not configure more than one SAP xserver resource on the same cluster because one SAP xserver serves multiple liveCache instances in the cluster. More than one SAP xserver resource that runs on the same cluster causes conflicts between the SAP xserver resources. These conflicts cause all SAP xserver resources to become unavailable. If you attempt to start the SAP xserver twice, you receive an error message that says `Address already in use`.

---

1. **Become superuser on one of the nodes in the cluster that will host the liveCache resource.**
2. **Copy the `lccluster` file to the same location as the `lcinit` file.**

```
cp /opt/SUNWsc1c/livecache/bin/lccluster \
/sapdb/LC-NAME/db/sap
```

LC-NAME    Uppercase name of liveCache database instance

3. Edit the `lccluster` file to substitute values for `put-LC_NAME-here` and `put-Confdir_list-here`.

---

**Note** – The `put-Confdir_list-here` value exists only in the Sun Cluster 3.1 version.

---

- a. Open the `lccluster` file.

```
vi /sapdb/LC-NAME/db/sap/lccluster \LC_NAME="put-LC_NAME-here" \
CONFDIR_LIST="put-Confdir_list-here"
```

---

**Note** – The `CONFDIR_LIST="put-Confdir_list-here` entry exists only in the Sun Cluster 3.1 version.

---

- b. Replace `put-LC_NAME-here` with the liveCache instance name. The liveCache instance name is the value you defined in the `Livecache_Name` extension property.

```
LC_NAME="liveCache-instance-name"
```

- c. Replace `put-Confdir_list-here` with the value of the `Confdir_list` extension property.

---

**Note** – This step is only for the Sun Cluster 3.1 version. Skip this step if you are running an earlier version of Sun Cluster.

---

```
CONFDIR_LIST="liveCache-software-directory"
```

Example:

If the liveCache instance name is LC1 and the liveCache software directory is `/sapdb`, edit the `lccluster` script as follows.

```
LC_NAME="LC1"
CONFDIR_LIST="/sapdb" [Sun Cluster 3.1 version only]
```

4. Add the `HASStoragePlus` resource to the liveCache resource group.

```
scrgadm -a -t SUNW.HASStoragePlus
scrgadm -a -j livecache-storage-resource -g livecache-resource-group \
-t SUNW.HASStoragePlus -x filesystemmountpoints=mountpoint,... \
-x globaldevicepaths=livecache-device-group
```

5. Enable the liveCache storage resource.

```
scswitch -e -j livecache-storage-resource
```

6. Register the resource type for liveCache database.

```
scrgadm -a -t SUNW.sap_livecache
```

7. Register the resource type for SAP xserver.

```
scrgadm -a -t SUNW.sap_xserver
```

8. Configure the SAP xserver as a scalable resource, completing the following substeps.

- a. Create a scalable resource group for SAP xserver. Configure SAP xserver to run on all the potential nodes that liveCache will run on.

---

**Note** – Configure SAP xserver so that SAP xserver starts on all nodes that the liveCache resources can fail over to. To implement this configuration, ensure that the *nodelist* parameter of the SAP xserver resource group contains all the nodes listed in the liveCache resource groups' *nodelist*. Also, the value of *desired primaries* and *maximum primaries* of the SAP xserver resource group must be equal to each other.

---

```
scrgadm -a -g xserver-resource-group \
-y Maximum_primaries=value \
-y Desired_primaries=value \
-h nodelist
```

- b. Create an SAP xserver resource in this scalable resource group.

```
scrgadm -a -j xserver-resource\
-g xserver-resource-group -t SUNW.sap_xserver
```

- c. Enable the scalable resource group that now includes the SAP xserver resource.

```
scswitch -Z -g xserver-resource-group
```

9. Register the liveCache resource as follows.

```
scrgadm -a -j livecache-resource -g livecache-resource-group \
-t SUNW.sap_livecache -x livecache_name=LC-NAME \
-y resource_dependencies=livecache-storage-resource
```

10. Set up a resource group dependency between SAP xserver and liveCache.

```
scrgadm -c -g livecache-resource-group \
-y rg_dependencies=xserver-resource-group
```

**11. Enable the liveCache failover resource group.**

```
scswitch -Z -g livecache-resource-group
```

**12. Are you running an APO application server on a node that liveCache can fail over to?**

- If no, this step completes this procedure.
- If yes, proceed to Step 9.

**13. Is the scalable APO application server resource group already in an RGOffload resource's rg\_to\_offload list?**

```
scrgadm -pvv | grep -i rg_to_offload | grep value:
```

- If yes, this step completes this procedure.
- If no, consider adding an RGOffload resource in the liveCache resource group.  
This configuration enables you to automatically shut down the APO application server if the liveCache resource fails over to a node on which the APO application server was running.

---

## Verifying the Sun Cluster HA for SAP liveCache Installation and Configuration

This section contains the procedure you need to verify that you installed and configured your data service correctly.

### ▼ How to Verify the Sun Cluster HA for SAP liveCache Installation and Configuration

Use this procedure to verify that you installed and configured Sun Cluster HA for SAP liveCache correctly. You need the information in the following table to understand the various states of the liveCache database.

**TABLE B-5** States of the liveCache database

| State   | Description                                     |
|---------|-------------------------------------------------|
| OFFLINE | liveCache is not running.                       |
| COLD    | liveCache is available for administrator tasks. |

**TABLE B-5** States of the liveCache database *(Continued)*

| State               | Description                                                                                                 |
|---------------------|-------------------------------------------------------------------------------------------------------------|
| WARM                | liveCache is online.                                                                                        |
| STOPPED INCORRECTLY | liveCache stopped incorrectly. This is also one of the interim states while liveCache starts or stops.      |
| ERROR               | Cannot determine the current state. This is also one of the interim states while liveCache starts or stops. |
| UNKNOWN             | This is one of the interim states while liveCache starts or stops.                                          |

1. **Log on to the node that hosts the resource group that contains the liveCache resource, and verify that the fault monitor functionality works correctly.**
  - a. **Terminate liveCache abnormally by stopping all liveCache processes.**  
Sun Cluster software restarts liveCache.  

```
ps -ef|grep sap|grep kernel
kill -9 livecache-processes
```
  - b. **Terminate liveCache by using the Stop liveCache button in LC10 or by running the `lcinit` command.**  
Sun Cluster software does not restart liveCache. However, the liveCache resource status message reflects that liveCache stopped outside of Sun Cluster software through the use of the Stop liveCache button in LC10 or the `lcinit` command. The state of the liveCache resource is UNKNOWN. When the user successfully restarts liveCache by using the Start liveCache button in LC10 or the `lcinit` command, the Sun Cluster HA for SAP liveCache Fault Monitor updates the resource state and status message to indicate that liveCache is running under the control of Sun Cluster software.
2. **Log on to SAP APO by using your SAP GUI with user DDIC, and verify that liveCache starts correctly by using transaction LC10.**
3. **As user root, switch the liveCache resource group to another node.**  

```
scswitch -z -g livecache-resource-group -h node2
```
4. **Repeat Step 1 through Step 3 for each potential node on which the liveCache resource can run.**
5. **Log on to the nodes that host the SAP xserver resource, and verify that the fault monitor functionality works correctly.**  
Terminate SAP xserver abnormally by stopping all SAP xserver processes.

```
ps -ef|grep xserver
kill -9 xserver-process
```

---

## Understanding Sun Cluster HA for SAP liveCache Fault Monitors

Use the information in this section to understand Sun Cluster HA for SAP liveCache Fault Monitors. This section describes the Sun Cluster HA for SAP liveCache Fault Monitors' probing algorithm or functionality, states the conditions, messages, and recovery actions associated with unsuccessful probing, and states the conditions and messages associated with successful probing.

### Extension Properties

See "Sun Cluster HA for SAP liveCache Extension Properties" on page 73 for the extension properties that the Sun Cluster HA for SAP liveCache Fault Monitors use.

### Monitor Check Method

A liveCache resource `Monitor_check` method checks whether SAP xserver is available on this node. If SAP xserver is not available on this node, this method returns an error and rejects the failover of liveCache to this node.

This method is needed to enforce the cross-resource group resource dependency between SAP xserver and liveCache.

### Probing Algorithm and Functionality

Sun Cluster HA for SAP liveCache has a fault monitor for each resource type.

- "SAP xserver Fault Monitor" on page 80 (`SUNW.sap_xserver`)
- "liveCache Fault Monitor" on page 81 (`SUNW.sap_livecache`)

### SAP xserver Fault Monitor

The SAP xserver parent process is under the control of process monitor `pmfadm`. If the parent process is stopped or killed, the process monitor contacts the SAP xserver Fault Monitor, and the SAP xserver Fault Monitor decides what action must be taken.



The SAP xserver Fault Monitor performs the following steps in a loop.

1. Sleeps for `Thorough_probe_interval`.
2. Uses the SAP utility `dbmcli` with `db_enum` to check SAP xserver availability.
  - If SAP xserver is unavailable, SAP xserver probe restarts or fails over the SAP xserver resource if it reaches the maximum number of restarts.
  - If any system error messages are logged in `syslog` during the checking process, the SAP xserver probe concludes that a partial failure has occurred. If the system error messages logged in `syslog` occur four times within the `probe_interval`, SAP xserver probe restarts SAP xserver.

## liveCache Fault Monitor

The liveCache probe checks for the presence of the liveCache parent process, the state of the liveCache database, and whether the user intentionally stopped liveCache outside of Sun Cluster software. If a user used the `Stop liveCache` button in LC10 or the `lcinit` command to stop liveCache outside of Sun Cluster software, the liveCache probe concludes that the user intentionally stopped liveCache outside of Sun Cluster software.

If the user intentionally stopped liveCache outside of Sun Cluster software by using the `Stop liveCache` button in LC10 or the `lcinit` command, the Sun Cluster HA for SAP liveCache Fault Monitor updates the resource state and status message to reflect this action, but it does not restart liveCache. When the user successfully restarts liveCache outside of Sun Cluster software by using the `Start liveCache` button in LC10 or the `lcinit` command, the Sun Cluster HA for SAP liveCache Fault Monitor updates the resource state and status message to indicate that liveCache is running under the control of Sun Cluster software, and Sun Cluster HA for SAP liveCache Fault Monitor takes appropriate action if it detects liveCache is `OFFLINE`.

If liveCache database state reports that liveCache is not running or that the liveCache parent process terminated, the Sun Cluster HA for SAP liveCache Fault Monitor restarts or fails over liveCache.

The Sun Cluster HA for SAP liveCache Fault Monitor performs the following steps in a loop. If any step returns `liveCache is offline`, the liveCache probe restarts or fails over liveCache.

1. Sleeps for `Thorough_probe_interval`.
2. Uses the `dbmcli` utility with `db_state` to check the liveCache database state.
3. If liveCache is online, liveCache probe checks the liveCache parent process.
  - If the parent process terminates, liveCache probe returns `liveCache is offline`.
  - If the parent process is online, liveCache probe returns `OK`.

4. If liveCache is not online, liveCache probe determines if the user stopped liveCache outside of Sun Cluster software by using the `Stop liveCache` button in LC10 or the `lcinit` command.
5. If the user stopped liveCache outside of Sun Cluster software by using the `Stop liveCache` button in LC10 or the `lcinit` command, returns OK.
6. If the user did not stop liveCache outside of Sun Cluster software by using the `Stop liveCache` button in LC10 or the `lcinit` command, checks SAP xserver availability.
  - If SAP xserver is unavailable, returns OK because the probe cannot restart liveCache if SAP xserver is unavailable.
  - If SAP xserver is available, returns `liveCache is offline`.
7. If any errors are reported from system function calls, returns `system failure`.

## Installing and Configuring Sun Cluster HA for Sybase ASE

---

This chapter provides instructions on how to configure and administer Sun Cluster HA for Sybase ASE on your Sun Cluster release notes documentation nodes.

This chapter contains the following procedures.

- “How to Prepare the Nodes” on page 85
- “How to Install the Sybase Software” on page 87
- “How to Verify the Sybase ASE Installation” on page 88
- “How to Configure Sybase ASE Database Access With Solstice DiskSuite/Solaris Volume Manager” on page 89
- “How to Configure Sybase ASE Database Access With VERITAS Volume Manager” on page 90
- “How to Create the Sybase ASE Database Environment” on page 91
- “How to Install Sun Cluster HA for Sybase ASE Packages” on page 93
- “How to Register and Configure Sun Cluster HA for Sybase ASE” on page 94
- “How to Verify the Sun Cluster HA for Sybase ASE Installation ” on page 97

You must configure Sun Cluster HA for Sybase ASE as a failover data service. See the Sun Cluster concepts documentation document and “Planning for Sun Cluster Data Services” in *Sun Cluster 3.0 12/01 Data Services Installation and Configuration Guide* for general information about data services, resource groups, resources, and other related topics.

---

## Installing and Configuring Sun Cluster HA for Sybase ASE

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

**TABLE C-1** Task Map: Installing and Configuring Sun Cluster HA for Sybase ASE

| Task                                                                                              | For Instructions, Go To                                                              |
|---------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------|
| Prepare to install Sun Cluster HA for Sybase ASE                                                  | "Preparing to Install Sun Cluster HA for Sybase ASE" on page 84                      |
| Install the Sybase ASE 12.0 software                                                              | "Installing the Sybase ASE 12.0 Software" on page 85                                 |
| Create the Sybase database environment                                                            | "Creating the Sybase ASE Database Environment" on page 89                            |
| Install the Sun Cluster HA for Sybase ASE package                                                 | "Installing the Sun Cluster HA for Sybase ASE Package" on page 92                    |
| Register Sun Cluster HA for Sybase ASE resource types and configure resource groups and resources | "Registering and Configuring Sun Cluster HA for Sybase ASE" on page 93               |
| Verify the Sun Cluster HA for Sybase ASE installation                                             | "Verifying the Sun Cluster HA for Sybase ASE Installation" on page 96                |
| Understand Sun Cluster HA for Sybase ASE logging and security issues                              | "Understanding Sun Cluster HA for Sybase ASE Logging and Security Issues" on page 98 |
| Configure Sun Cluster HA for Sybase ASE extension properties                                      | "Configuring Sun Cluster HA for Sybase ASE Extension Properties" on page 99          |
| View fault monitor information                                                                    | "Sun Cluster HA for Sybase ASE Fault Monitor" on page 101                            |

---

## Preparing to Install Sun Cluster HA for Sybase ASE

To prepare your nodes for the Sun Cluster release notes documentation HA for Sybase Adaptive Server 12.0 installation, select an installation location for the following files.

- **Sybase ASE application files** – These files include Sybase ASE binaries and libraries. You can install these files on either the local file system or the cluster file system.  
See the Sun Cluster data services collection for the advantages and disadvantages of placing the Sybase ASE binaries on the local file system as opposed to the cluster file system.
- **Sybase ASE configuration files** – These files include the `interfaces` file, `config` file, and `environment` file. You can install these files on the local file system (with links), the highly available local file system, or on the cluster file system.

- **Database data files** – These files include Sybase device files. You must install these files on the highly available local file system or the cluster file system as either raw devices or regular files.

---

## Installing the Sybase ASE 12.0 Software

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

- Prepare the nodes.
- Install the Sybase ASE software.
- Verify the Sybase ASE installation.

---

**Note** – Before you configure Sun Cluster HA for Sybase ASE, use the procedures that the Sun Cluster data services collection describes to configure the Sybase ASE software on each node.

---

### ▼ How to Prepare the Nodes

This procedure describes how to prepare the cluster nodes for Sybase ASE software installation.



---

**Caution** – Perform all of the steps in this procedure on all of the nodes. If you do not perform all of the steps on all of the nodes, the Sybase ASE installation will be incomplete, and Sun Cluster HA for Sybase ASE will fail during startup.

---

---

**Note** – Consult the Sybase ASE documentation before you perform this procedure.

---

1. **Become superuser on all of the nodes.**
2. **Configure the `/etc/nsswitch.conf` file as follows so that Sun Cluster HA for Sybase ASE starts and stops correctly if a switchover or failover occurs.**

On each node that can master the logical host that runs Sun Cluster HA for Sybase ASE, include one of the following entries for `group` in the `/etc/nsswitch.conf` file.

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

```
group: file [NOTFOUND=return] nisplus
```

Sun Cluster HA for Sybase ASE 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. Adding one of the preceding entries for `group` ensures that the `su(1M)` command does not refer to the NIS/NIS+ name services if the network information name service is unavailable.

### 3. Configure the cluster file system for Sun Cluster HA for Sybase ASE.

If raw devices contain the databases, configure the global devices for raw-device access. See the Sun Cluster data services collection for information on how to configure global devices.

If you use the Solstice DiskSuite/Solaris Volume Manager volume manager, configure the Sybase ASE software to use UNIX file system (UFS) logging on mirrored meta devices or raw-mirrored meta devices. See the Solstice DiskSuite/Solaris Volume Manager documentation for information on how to configure raw-mirrored metadevices.

### 4. Prepare the `SYBASE_HOME` directory on a local or multihost disk.

---

**Note** – If you install the Sybase ASE binaries on a local disk, use a separate disk if possible. Installing the Sybase ASE binaries on a separate disk prevents the binaries from overwrites during operating environment reinstallation.

---

### 5. On each node, create an entry for the database administrator (DBA) group in the `/etc/group` file, and add potential users to the group.

Verify that the `root` and `sybase` users are members of the `dba` group, and add entries as necessary for other DBA users. Ensure that group IDs are the same on all of the nodes that run Sun Cluster HA for Sybase ASE, as the following example illustrates.

```
dba:*:520:root,sybase
```

You can create group entries in a network name service. If you do so, also add your entries to the local `/etc/group` file to eliminate dependency on the network name service.

### 6. On each node, create an entry for the Sybase system administrator.

The following command updates the `/etc/passwd` and `/etc/shadow` files with an entry for the Sybase system administrator.

```
useradd -u 120 -g dba -d /Sybase-home sybase
```

Ensure that the `sybase` user entry is the same on all of the nodes that run Sun Cluster HA for Sybase ASE.

## ▼ How to Install the Sybase Software

Perform the following steps to install the Sybase ASE software.

### 1. Become superuser on a cluster member.

### 2. Note the Sybase ASE installation requirements.

You can install Sybase ASE binaries on one of the following locations.

- Local disks of the cluster nodes
- Highly available local file system
- Cluster file system

---

**Note** – Before you install the Sybase ASE software on the cluster file system, start the Sun Cluster release notes documentation software and become the owner of the disk device group.

---

See “Preparing to Install Sun Cluster HA for Sybase ASE” on page 84 for more information about installation locations.

### 3. Create a failover resource group to hold the network and application resources.

```
scrgadm -a -g resource-group [-h nodelist]
```

`-g resource-group` Specifies the name of the resource group. This name can be your choice but must be unique for resource groups within the cluster.

`-h nodelist` Specifies an optional, comma-separated list of physical node names or IDs that identify potential masters. The order here determines the order in which the Resource Group Manager (RGM) considers primary nodes during failover.

---

**Note** – Use the `-h` option to specify the order of the node list. If all of the nodes in the cluster are potential masters, you do not need to use the `-h` option.

---

### 4. Verify that you have added all of the network resources that Sun Cluster HA for Sybase ASE uses to either the `/etc/inet/hosts` file or to your name service (NIS, NIS+) database.

### 5. Add a network resource (logical hostname or shared address) to the failover resource group.

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

- l *logical-hostname* Specifies a network resource. The network resource is the logical hostname or shared address (IP address) that clients use to access Sun Cluster HA for Sybase ASE.
- n *netiflist* Specifies an optional, comma-separated list that identifies the NAFO groups on each node. All of the nodes that are in the resource group's *nodelist* must be represented in the *netiflist*. If you do not specify this option, the *scrgadm(1M)* command attempts to discover a net adapter on the subnet that the hostname list identifies for each node that is in *nodelist*. For example, -n *nafo0@nodename, nafo0@nodename2*.

**6. Run the `scswitch(1M)` command to complete the following tasks..**

- Enable the resource and fault monitoring.
- Move the resource group into a managed state.
- Bring the resource group online.

```
scswitch -z -g resource-group
```

**7. On the node mastering the resource group that you just created, login as *sybase*.**

The installation of the Sybase binaries must be performed on the node where the corresponding logical host is running.

**8. Install the Sybase ASE software.**

Regardless of where you install the Sybase ASE software, modify each node's `/etc/system` files as you would in standard Sybase ASE installation procedures. For instructions on how to install the Sybase ASE software, refer to the Sybase installation and configuration guides.

---

**Note** – For every Sybase server, enter the hostname that is associated with a network resource when asked to specify the hostname.

---

## Where to Go From Here

After you install the Sybase ASE software, go to “How to Configure Sybase ASE Database Access With Solstice DiskSuite/Solaris Volume Manager” on page 89 if you use the Solstice DiskSuite/Solaris Volume Manager volume manager. Go to “How to Configure Sybase ASE Database Access With VERITAS Volume Manager” on page 90 if you use the VERITAS Volume Manager (VxVM).

## ▼ How to Verify the Sybase ASE Installation

Perform the following steps to verify the Sybase ASE software installation.



1. Verify that the *sybase* user and the *dba* group own the `$SYBASE_HOME` directory and `$SYBASE_HOME` children directories.
2. Run the `scstat(1M)` command to verify that the Sun Cluster release notes documentation software functions correctly.

---

## Creating the Sybase ASE Database Environment

The procedures in this section enable you to complete the following tasks.

- Configure Sybase ASE database access with Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager.
- Create the Sybase ASE database environment.

### ▼ How to Configure Sybase ASE Database Access With Solstice DiskSuite/Solaris Volume Manager

If you use the Solstice DiskSuite/Solaris Volume Manager volume manager, perform the following steps to configure Sybase ASE database access with the Solstice DiskSuite/Solaris Volume Manager volume manager.

1. **Configure the disk devices for the Solstice DiskSuite/Solaris Volume Manager software to use.**

See the Sun Cluster software installation documentation for information on how to configure Solstice DiskSuite/Solaris Volume Manager.

2. **If you use raw devices to contain the databases, run the following commands to change each raw-mirrored metadevice's owner, group, and mode.**

If you do not use raw devices, do not perform this step.

- a. **If you create raw devices, run the following commands for each device on each node that can master the Sybase ASE resource group.**

```
chown sybase /dev/md/metaset/rdisk/dn
chgrp dba /dev/md/metaset/rdisk/dn
chmod 600 /dev/md/metaset/rdisk/dn
```

*metaset* Specifies the name of the diskset.

`/rdsk/dn` Specifies the name of the raw disk device within the *metaset* diskset.

**b. Verify that the changes are effective.**

```
ls -lL /dev/md/metaset/rdsk/dn
```

## ▼ How to Configure Sybase ASE Database Access With VERITAS Volume Manager

If you use VERITAS Volume Manager software, perform the following steps to configure Sybase ASE database access with the VERITAS Volume Manager software.

**1. Configure the disk devices for the VERITAS Volume Manager software to use.**

See the Sun Cluster software installation documentation for information on how to configure VERITAS Volume Manager.

**2. If you use raw devices to contain the databases, run the following commands on the current disk-group primary to change each device's owner, group, and mode.**

If you do not use raw devices, do not perform this step.

**a. If you create raw devices, run the following command for each raw device.**

```
vxedit -g diskgroup set user=sybase group=dba mode=0600 volume
```

`-g resource-group` Specifies the name of the resource group. This name can be your choice but must be unique for resource groups within the cluster.

`-h nodelist` Specifies an optional comma-separated list of physical node names or IDs that identify potential masters. The order here determines the order in which the nodes are considered as primary during failover.

**b. Verify that the changes are effective.**

```
ls -lL /dev/vx/rdsk/diskgroup/volume
```

**c. Reregister the disk device group with the cluster to keep the VERITAS Volume Manager namespace consistent throughout the cluster.**

```
scconf -c -D name=diskgroup
```

## ▼ How to Create the Sybase ASE Database Environment

Before you perform this procedure, ensure that you have completed the following tasks.

- Establish a highly available IP address and name, that is, a network resource that operates at installation time.
- Locate device paths for all of the Sybase ASE devices—including the master device and system devices—in the highly available local file system or cluster file system. Configure device paths as one of the following file types.
  - regular files
  - raw devices
  - files that the Solstice DiskSuite/Solaris Volume Manager software or the VERITAS Volume Manager software manage
- Locate the Sybase ASE server logs in either the cluster file system or the local file system.
- The Sybase ASE 12.0 environment consists of the data server, backup server, monitor server, text server, and XP server. The data server is the only server that you must configure—you can choose whether to configure all of the other servers.
- The entire cluster must contain only one copy of the `interfaces` file. The `§SYBASE` directory contains the `interfaces` file. If you plan to maintain per-node file copies, ensure the file contents are identical.

All of the clients that connect to Sybase ASE servers connect with Sybase OpenClient libraries and utilities. When you configure the Sybase ASE software, in the `interfaces` file, enter information about the network resource and various ports. All of the clients use this connection information to connect to the Sybase ASE servers.

Perform the following steps to create the Sybase ASE database environment.

**1. Run the GUI-based utility `srvbuild` to create the Sybase ASE database.**

The `§SYBASE/ASE_12-0/bin` directory contains this utility. See the Sybase ASE document entitled “Installing Sybase Adaptive Server Enterprise on Sun Solaris 2.x (SPARC).”

**2. To verify successful database installation, ensure that all of the servers start correctly.**

Run the `ps(1)` command to verify the operation of all of the servers. Sybase ASE server logs indicate any errors that have occurred.

**3. Set the password for the Sybase ASE system administrator account.**

See the *Sybase Adaptive Server Enterprise System Administration Guide* for details on changing the `sa` login password.

**4. Create a new Sybase ASE account for fault monitoring.**

This account enables the fault monitor to perform the following tasks.

- Support queries to system tables.
- Create and update user tables.

---

**Note** – Do not use the `sa` account for these purposes.

---

See “Sun Cluster HA for Sybase ASE Fault Monitor” on page 101 for more information.

**5. Update the stop file with the `sa` password.**

Because the stop file contains the `sa` password, protect the file with the appropriate permissions, and place the file in a directory that the system administrator chooses. Enable only the `sybase` user to read, write, and execute the stop file.

See “Important Security Issues” on page 98 for more information about the stop file.

## Where to Go From Here

After you create the Sybase ASE database environment, go to “How to Install Sun Cluster HA for Sybase ASE Packages” on page 93.

---

# Installing the Sun Cluster HA for Sybase ASE Package

You can use the `scinstall(1M)` utility to install `SUNWscsyb`, the Sun Cluster HA for Sybase ASE package, on a cluster. Do not use the `-s` option to non-interactive `scinstall` to install all of the data service packages.

If you installed the `SUNWscsyb` data service package as part of your initial Sun Cluster release notes documentation installation, proceed to “Registering and Configuring Sun Cluster HA for Sybase ASE” on page 93. Otherwise, use the following procedure to install the `SUNWscsyb` package.

## ▼ How to Install Sun Cluster HA for Sybase ASE Packages

You need the Sun Cluster 3.1 Agents CD-ROM to complete this procedure. Perform this procedure on all of the cluster nodes that run the Sun Cluster HA for Sybase ASE package.

1. **Load the Sun Cluster 3.1 Agents CD-ROM into the CD-ROM drive.**
2. **Run the `scinstall` utility with no options.**  
This step starts the `scinstall` utility in interactive mode.
3. **Choose the menu option, Add Support for New Data Service to This Cluster Node.**  
The `scinstall` utility prompts you for additional information.
4. **Provide the path to the Sun Cluster 3.1 Agents CD-ROM.**  
The utility refers to the CD as the “data services cd.”
5. **Specify the data service to install.**  
The `scinstall` utility lists the data service that you selected and asks you to confirm your choice.
6. **Exit the `scinstall` utility.**
7. **Unload the CD from the drive.**

### Where to Go From Here

When you finish the Sun Cluster HA for Sybase ASE package installation, go to “How to Register and Configure Sun Cluster HA for Sybase ASE” on page 94.

---

## Registering and Configuring Sun Cluster HA for Sybase ASE

Use the procedures in this section to register and configure the Sun Cluster HA for Sybase ASE data service. Register and configure Sun Cluster HA for Sybase ASE as a failover data service.

## ▼ How to Register and Configure Sun Cluster HA for Sybase ASE

This procedure describes how to use the `scrgadm(1M)` command to register and configure Sun Cluster HA for Sybase ASE.

This procedure includes creating the `HASStoragePlus` resource type. This resource type synchronizes actions between `HASStorage` and Sun Cluster HA for Sybase ASE and enables you to use a highly available local file system. Sun Cluster HA for Sybase ASE is disk-intensive, and therefore you should configure the `HASStoragePlus` resource type.

See the `SUNW.HASStoragePlus(5)` man page and “Relationship Between Resource Groups and Disk Device Groups” in *Sun Cluster 3.0 12/01 Data Services Installation and Configuration Guide* for more information about the `HASStoragePlus` resource type.

---

**Note** – Other options also enable you to register and configure the data service. See “Tools for Data Service Resource Administration” in *Sun Cluster 3.0 12/01 Data Services Installation and Configuration Guide* for details about these options.

---

To perform this procedure, you must have the following information.

- The names of the cluster nodes that master the data service.
- The network resource that clients use to access the data service. You typically configure the IP address when you install the cluster. See the sections in the Sun Cluster software installation documentation on planning the Sun Cluster environment and on how to install the Solaris operating environment for details.
- The path to the Sybase ASE application installation.

---

**Note** – Perform the following steps on one cluster member.

---

1. **Become superuser on a cluster member.**
2. **Run the `scrgadm` command to register resource types for Sun Cluster HA for Sybase ASE.**

```
scrgadm -a -t SUNW.sybase
```

|                             |                                                                            |
|-----------------------------|----------------------------------------------------------------------------|
| -a                          | Adds the resource type for the data service.                               |
| -t <code>SUNW.sybase</code> | Specifies the resource type name that is predefined for your data service. |

3. **Register the `HASStoragePlus` resource type with the cluster.**

```
scrgadm -a -t SUNW.HAStoragePlus
```

4. Create the resource `sybase-hastp-rs` of type `HAStoragePlus`.

```
scrgadm -a -j sybase-hastp-rs -g sybase-rg -t SUNW.HAStoragePlus \
-x GlobalDevicePaths=sybase-set1,/dev/global/dsk/d1 \
-x FilesystemMountPoints=/global/sybase-inst \
-x AffinityOn=TRUE
```

---

**Note** – AffinityOn must be set to TRUE and the local file system must reside on global disk groups to be failover.

---

5. Run the `scrgadm` command to complete the following tasks and bring the resource group `sybase-rg` online on a cluster node.

- Move the resource group into a managed state.
- Bring the resource group online

This node will be made the primary for device group `sybase-set1` and raw device `/dev/global/dsk/d1`. Device groups associated with filesystems such as `/global/sybase-inst` will also be made primaries on this node..

```
scrgadm -Z -g sybase-rg
```

6. Create Sybase ASE application resources in the failover resource group.

```
scrgadm -a -j resource -g resource-group \
-t SUNW.sybase \
-x Environment_File=environment-file-path \
-x Adaptive_Server_Name=adaptive-server-name \
-x Backup_Server_Name=backup-server-name \
-x Text_Server_Name=text-server-name \
-x Monitor_Server_Name=monitor-server-name \
-x Adaptive_Server_Log_File=log-file-path \
-x Stop_File=stop-file-path \
-x Connect_string=user/passwd \
-y resource_dependencies=storageplus-resource
```

-j *resource*

Specifies the resource name to add.

-g *resource-group*

Specifies the resource group name into which the RGM places the resources.

-t `SUNW.sybase`

Specifies the resource type to add.

-x `Environment_File=environment-file`

Sets the name of the environment file.

-x `Adaptive_Server_Name=adaptive-server-name`

Sets the name of the adaptive server.

- x Backup\_Server\_Name=*backup-server-name*  
Sets the name of the backup server.
- x Text\_Server\_Name=*text-server-name*  
Sets the name of the text server.
- x Monitor\_Server\_Name=*monitor-server-name*  
Sets the name of the monitor server.
- x Adaptive\_Server\_Log\_File=*log-file-path*  
Sets the path to the log file for the adaptive server.
- x Stop\_File=*stop-file-path*  
Sets the path to the stop file.
- x Connect\_string=*user/passwd*  
Specifies the user name and password that the fault monitor uses to connect to the database.

You do not have to specify extension properties that have default values. See “Configuring Sun Cluster HA for Sybase ASE Extension Properties” on page 99 for more information.

## 7. Enable the resource and fault monitoring.

---

**Note** – Sybase start logs print to the console when the Sybase servers start. If you do not want these messages to print to the console, update the appropriate RUN files to redirect these messages to another file.

---

```
scswitch -Z -g resource-group
```

## Where to Go From Here

After you register and configure Sun Cluster HA for Sybase ASE, go to “How to Verify the Sun Cluster HA for Sybase ASE Installation ” on page 97.

---

# Verifying the Sun Cluster HA for Sybase ASE Installation

Perform the following verification tests to ensure that you have correctly installed and configured Sun Cluster HA for Sybase ASE.



These sanity checks ensure that all of the nodes that run Sun Cluster HA for Sybase ASE can start the Sybase ASE data server. These checks also ensure that other nodes in the configuration can access the Sybase ASE data server. Perform these sanity checks to isolate any problems with starting the Sybase ASE software from Sun Cluster HA for Sybase ASE.

## ▼ How to Verify the Sun Cluster HA for Sybase ASE Installation

1. **Log in to the node that masters the Sybase ASE resource group.**

2. **Set the Sybase ASE environment variables.**

The environment variables are the variables that you specify with the `Environment_file` extension property. You typically name this file `SYBASE.sh` or `SYBASE.csh`.

3. **Verify that the Sun Cluster HA for Sybase ASE resource is online.**

```
scstat -g
```

4. **Inspect the Sybase ASE logs to determine the cause of any errors that have occurred.**

5. **Confirm that you can connect to the data server and execute the following test command.**

```
isql -S adaptive-server -U sa
```

```
isql> sp_help
```

```
isql> go
```

```
isql> quit
```

6. **Kill the process for the Sybase ASE data server.**

The Sun Cluster release notes documentation software restarts the process.

7. **Switch the resource group that contains the Sybase ASE resource to another cluster member.**

```
scswitch -z -g resource-group -h node
```

8. **Log in to the node that now contains the resource group.**

9. **Repeat Step 3 and Step 5.**

---

**Note** – Sybase ASE client connections cannot survive a Sun Cluster HA for Sybase ASE switchover. If a switchover occurs, the existing client connections to Sybase ASE terminate, and clients must reestablish their connections. After a switchover, the time that is required to replay the Sybase ASE transaction log determines Sun Cluster HA for Sybase ASE recovery time.

---

---

## Understanding Sun Cluster HA for Sybase ASE Logging and Security Issues

The following sections contain information about Sun Cluster HA for Sybase ASE logging and security issues.

### Sun Cluster HA for Sybase ASE Logging

Sun Cluster HA for Sybase ASE logs messages to the file `message_log` in the `/opt/SUNWscsyb/1` directory. Although this file cannot exceed 512 Kbytes, Sun Cluster HA for Sybase ASE does not delete old log files. The number of log files, therefore, can grow to a large number.

Sun Cluster HA for Sybase ASE writes all of the error messages in the `syslog` file. Sun Cluster HA for Sybase ASE also logs fault monitor history to the file `restart_history` in the `log` directory. These files can also grow to a large number.

As part of your regular file maintenance, check the following log files and remove files that you no longer need.

- `syslog`
- `message_log`
- `restart_history`

### Important Security Issues

Sun Cluster HA for Sybase ASE requires that you embed the system administrator's password in a stop file. The `/opt/SUNWscsyb/bin` directory contains the template for the stop file, `sybase_stop_servers`. Sun Cluster HA for Sybase ASE uses this file to log in to the Sybase ASE environment and to stop the Sybase ASE servers. Enable the `sybase` user to execute the stop file, but protect the file from general access. Give read, write, and execute privileges to only the following users.

- *sybase* user
- *sybase* group

---

## Configuring Sun Cluster HA for Sybase ASE Extension Properties

This section describes how to configure Sun Cluster HA for Sybase ASE extension properties. Typically, you use the command line `scrgadm -x parameter=value` to configure extension properties when you create the Sybase ASE resources. You can also use the procedures that “Administering Data Service Resources” in *Sun Cluster 3.0 12/01 Data Services Installation and Configuration Guide* describes to configure them later.

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

Table C-2 describes the extension properties that you can set for the Sybase ASE server resource. You can update some extension properties dynamically. You can update others, however, only when you create or disable a resource. The Tunable entries indicate when you can update each property.

**TABLE C-2** Sun Cluster HA for Sybase ASE Extension Properties

| Name/Data Type                   | Description                                                                                                                                                                                                                                                         |
|----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>Environment_File</code>    | <p>File that contains all of the Sybase ASE environment variables. This file is automatically created in the Sybase home directory.</p> <p><b>Default:</b> None</p> <p><b>Range:</b> Minimum=1</p> <p><b>Tunable:</b> When disabled</p>                             |
| <code>Adaptive_ServerName</code> | <p>The name of the data server. Sun Cluster HA for Sybase ASE uses this property to locate the RUN server in the <code>\$SYBASE/\$ASE/install</code> directory.</p> <p><b>Default:</b> None</p> <p><b>Range:</b> Minimum=1</p> <p><b>Tunable:</b> When disabled</p> |

**TABLE C-2** Sun Cluster HA for Sybase ASE Extension Properties (Continued)

| Name/Data Type       | Description                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Backup_Server_Name   | <p>The name of the backup server. Sun Cluster HA for Sybase ASE uses this property to locate the RUN server in the \$SYBASE/\$ASE/install directory. If you do not set this property, Sun Cluster HA for Sybase ASE will not manage the server.</p> <p><b>Default:</b> Null</p> <p><b>Range:</b> None</p> <p><b>Tunable:</b> When disabled</p>                                 |
| Monitor_Server_Name  | <p>The name of the monitor server. Sun Cluster HA for Sybase ASE uses this property to locate the RUN server in the \$SYBASE/\$ASE/install directory. If you do not set this property, Sun Cluster HA for Sybase ASE will not manage the server.</p> <p><b>Default:</b> Null</p> <p><b>Range:</b> None</p> <p><b>Tunable:</b> When disabled</p>                                |
| Text_Server_Name     | <p>The name of the text server. The Sun Cluster HA for Sybase ASE data service uses this property to locate the RUN server in the \$SYBASE/\$ASE/install directory. If you do not set this property, the Sun Cluster HA for Sybase ASE data service will not manage the server.</p> <p><b>Default:</b> Null</p> <p><b>Range:</b> None</p> <p><b>Tunable:</b> When disabled</p> |
| Adaptive_Server_File | <p>The path to the log file for the adaptive server. Sun Cluster HA for Sybase ASE continually reads this property for error monitoring.</p> <p><b>Default:</b> None</p> <p><b>Range:</b> Minimum=1</p> <p><b>Tunable:</b> When disabled</p>                                                                                                                                   |
| Stop_File            | <p>Sun Cluster HA for Sybase ASE uses this property during server stoppages. This property contains the sa password. Protect this property from general access.</p> <p><b>Default:</b> None</p> <p><b>Range:</b> Minimum=1</p> <p><b>Tunable:</b> When disabled</p>                                                                                                            |

**TABLE C-2** Sun Cluster HA for Sybase ASE Extension Properties (Continued)

| Name/Data Type  | Description                                                                                                                                                                                       |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Probe_timeout   | Time-out value for the fault monitor probe.<br><b>Default:</b> 30 seconds<br><b>Range:</b> 1 – 99999 seconds<br><b>Tunable:</b> Any time                                                          |
| Debug_level     | Debug level for writing to the Sun Cluster HA for Sybase ASE log.<br><b>Default:</b> 0<br><b>Range:</b> 0 – 15<br><b>Tunable:</b> Any time                                                        |
| Connect_string  | String of format <i>user/password</i> . Sun Cluster HA for Sybase ASE uses this property for database probes.<br><b>Default:</b> None<br><b>Range:</b> Minimum=1<br><b>Tunable:</b> When disabled |
| Connect_cycle   | Number of fault monitor probe cycles before Sun Cluster HA for Sybase ASE establishes a new connection.<br><b>Default:</b> 5<br><b>Range:</b> 1 – 100<br><b>Tunable:</b> Any time                 |
| Wait_for_online | Whether the start method waits for the database to come online before exiting.<br><b>Default:</b> FALSE<br><b>Range:</b> TRUE – FALSE<br><b>Tunable:</b> Any time                                 |

---

## Sun Cluster HA for Sybase ASE Fault Monitor

The Sun Cluster HA for Sybase ASE fault monitor queries the Sybase ASE server to determine server health.

---

**Note** – The Sun Cluster HA for Sybase ASE fault monitor only monitors the Adaptive server. The fault monitor does not monitor auxiliary servers.

---

The fault monitor consists of the following processes.

- a main fault monitor process
- a database-client fault probe

The following sections describe the Sun Cluster HA for Sybase ASE fault monitor processes and the extension properties that the fault monitor uses.

## Main Fault Monitor Process

The fault monitor process diagnoses errors and checks statistics. The monitor labels an operation successful if the following conditions occur.

- The database is online.
- The activity check returns no errors.
- The test transaction returns no errors.

If an operation fails, the main process checks the action table for an action to perform and then performs the predetermined action. If an operation fails, the main process can perform the following actions, which execute external programs as separate processes in the background.

1. Restarts the resource on the current node.
2. Restarts the resource group on the current node.
3. Fails over the resource group to the next node on the resource group's nodelist.

The server fault monitor also scans the `Adaptive_Server_Log` file and acts to correct any errors that the scan identifies.

## Database-Client Fault Probe

The database-client fault probe performs activity checks and test transactions. The extension property `Connect_string` specifies an account that performs all of the database operations. The extension property `Probe_timeout` sets the time-out value that the probe uses to determine the time that has elapsed in a successful database probe.

## Extension Properties

The fault monitor uses the following extension properties.

- Thorough\_probe\_interval
- Retry\_count
- Retry\_interval
- Probe\_timeout
- Connect\_string
- Connect\_cycle
- Adaptive\_Server\_Log

See “Configuring Sun Cluster HA for Sybase ASE Extension Properties” on page 99 for more information about these extension properties.





## RSM Phase II: RSMRDT Driver Installation

---

This appendix describes the prerequisites and procedures for installation of the Remote Shared Memory Reliable Datagram Transport (RSMRDT) driver. This appendix includes the following sections:

---

**Note** – The RSMRDT driver should not be installed until RSM with 9iRAC is supported. Contact your Sun service provider for configuration support information.

---

- “Overview of the RSMRDT Driver” on page 105
- “Restrictions” on page 106
- “How to Install the SUNWscrDt Package” on page 106
- “How to Uninstall the SUNWscrDt Package” on page 106
- “How to Unload the RSMRDT Driver Manually” on page 107

---

### Overview of the RSMRDT Driver

Remote Shared Memory (RSM) is an interface on top of a memory-based interconnect. RSM provides highly reliable remote memory operations and synchronous detection of communication failure through barrier calls. RSMRDT consists of a driver that is built on top of RSMPI and a library that exports the RSMRDT-API interface. RSMRDT is dependent on Sun Cluster software and RSM. The primary goal of the driver is to provide enhanced Oracle™ Parallel Server (OPS) performance. A secondary goal is to enhance load-balancing and high-availability (HA) functions by providing them directly inside the driver, making them available to the clients transparently.

## Installing the RSMRDT Driver

The RSMRDT driver and library are installed with the `SUNWscrdt` package. You must successfully install Sun Cluster software, the RSM package, `SUNWrsmo`, and `SUNWrsmx` before beginning RSMRDT installation.

## Restrictions

Use of the RSMRDT Driver is restricted to customers running an Oracle9i release 2 SCI configuration with RSM enabled. Refer to Oracle9i release 2 user documentation for detailed installation and configuration instructions. The `SUNWscrdt` package (RSMRDT driver package) depends on the following packages:

- `SUNWrsmo` RSMPI Operations Registration Module
- `SUNWrsmox` RSMPI Operations Registration Module (64-bit)

The `SUNWscrdt` package also has a functional dependency on the following RSM packages:

- `SUNWrsm` Remote Shared Memory
- `SUNWrsmx` Remote Shared Memory (64-bit)

## ▼ How to Install the `SUNWscrdt` Package

1. Verify that `SUNWrsmo` and `SUNWrsmx` are installed before completing this procedure.
2. Become superuser on the node to which you want to install the `SUNWscrdt` package.

---

**Note** – You must repeat this procedure for each node in the cluster.

---

3. Install the `SUNWscrdt` package.

```
pkgadd -d pathname SUNWscrdt
```

*pathname* Specifies the path name of the directory that contains `SUNWscrdt`

## ▼ How to Uninstall the `SUNWscrdt` Package

1. Verify that no applications are using the RSMRDT driver before performing this procedure.

2. Become superuser on the node to which you want to uninstall the `SUNWscrtd` package.

---

**Note** – You must repeat this procedure for each node in the cluster.

---

3. Uninstall the `SUNWscrtd` package.

```
pkgrm SUNWscrtd
```

## ▼ How to Unload the RSMRDT Driver Manually

If the driver remains loaded in memory after completing “How to Uninstall the `SUNWscrtd` Package” on page 106, perform the following procedure to unload the driver manually.

1. Start the `adb` tool.

```
adb -kw
```

2. Set the kernel variable `clif_rsmrdt_modunload_ok` to 1.

```
physmem ####
clif_rsmrdt_modunload_ok/W 1
```

3. Exit `adb` by pressing Control-D.

4. Find the `clif_rsmrdt` and `rsmrdt` module IDs.

```
modinfo | grep rdt
```

5. Unload the `clif_rsmrdt` module.

```
modunload -i clif_rsmrdt_id
```

---

**Note** – You must unload the `clif_rsmrdt` module before unloading `rsmrdt`. If `modunload` fails, applications are probably still using the driver. Terminate the applications before running `modunload` again.

---

*clif\_rsmrdt\_id* Specifies the numeric ID for the module being unloaded.

6. Unload the `rsmrdt` module.

```
modunload -i rsmrdt_id
```

*rsmrdt\_id* Specifies the numeric ID for the module being unloaded.

## 7. Verify that the module was unloaded successfully.

```
modinfo | grep rdt
```

### Example—Unloading the RSMRDT Driver

The following example shows the console output after the RSMRDT driver is manually unloaded.

```
adb -kw
physmem fc54
clifrsrmt_modunload_ok/W 1
clifrsrmt_modunload_ok: 0x0 = 0x1
^D
modinfo | grep rsm
 88 f064a5cb 974 - 1 rsmops (RSMOPS module 1.1)
 93 f08e07d4 b95 - 1 clif_rsrmt (CLUSTER-RSMRDT Interface module)
 94 f0d3d000 13db0 194 1 rsrmt (Reliable Datagram Transport dri)
modunload -i 93
modunload -i 94
modinfo | grep rsm
 88 f064a5cb 974 - 1 rsmops (RSMOPS module 1.1)
#
```

# Installing and Configuring Sun Cluster HA for SAP

---

This appendix contains the procedures on how to install and configure Sun Cluster HA for SAP.

This appendix contains the following procedures.

- “How to Upgrade a Resource Type or Convert a Failover Application Resource to a Scalable Application Resource” on page 118
- “How to Prepare the Nodes” on page 119
- “How to Install SAP and the Database” on page 120
- “How to Install an SAP Scalable Application Server” on page 120
- “How to Verify an SAP Scalable Application Server” on page 128
- “How to Enable Failover SAP Instances to Run in a Cluster” on page 123
- “How to Verify SAP and the Database Installation with Central Instance” on page 125
- “How to Verify an SAP Failover Application Server” on page 126
- “How to Install the Sun Cluster HA for SAP Packages” on page 128
- “How to Register and Configure Sun Cluster HA for SAP with Central Instance” on page 136
- “How to Register and Configure Sun Cluster HA for SAP as a Failover Data Service” on page 137
- “How to Register and Configure Sun Cluster HA for SAP as a Scalable Data Service” on page 138
- “How to Set Up a Lock File for Central Instance or the Failover Application Server” on page 140
- “How to Set Up a Lock File for Scalable Application Server” on page 141
- “How to Verify Sun Cluster HA for SAP Installation and Configuration and Central Instance” on page 142
- “How to Verify the Installation and Configuration of Sun Cluster HA for SAP as a Failover Data Service” on page 143
- “How to Verify Sun Cluster HA for SAP Installation and Configuration of as a Scalable Data Service” on page 143

---

## Sun Cluster HA for SAP Overview

Use the information in this section to understand how Sun Cluster HA for SAP makes SAP highly available.

For conceptual information on failover and scalable services, see the *Sun Cluster concepts documentation*.

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.

**TABLE E-1** Protection of SAP Components

| SAP Component          | Protected by                                                                                             |
|------------------------|----------------------------------------------------------------------------------------------------------|
| SAP database           | Sun Cluster HA for Oracle<br>Use Oracle as your database.                                                |
| SAP central instance   | Sun Cluster HA for SAP<br>The resource type is <code>SUNW.sap_ci</code> or <code>SUNW.sap_ci_v2</code> . |
| SAP application server | Sun Cluster HA for SAP<br>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 software installation documentation* 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

Table E-2 lists the tasks for installing and configuring Sun Cluster HA for SAP. Perform these tasks in the order that they are listed.

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

| Task                                                           | For Instructions, Go To                                                                                                                                                   |
|----------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Plan the SAP installation                                      | Chapter 1 of <i>Sun Cluster 3.x Release Notes Supplement</i><br>"Planning the Sun Cluster HA for SAP Installation and Configuration" on page 112                          |
| 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 118                                                |
| Prepare the nodes and disks                                    | "How to Prepare the Nodes" on page 119                                                                                                                                    |
| Install SAP, SAP failover application server, and the database | "How to Install SAP and the Database" on page 120<br>"How to Enable Failover SAP Instances to Run in a Cluster" on page 123                                               |
| Configure the Sun Cluster HA for DBMS                          | "Configuring Sun Cluster HA for DBMS" on page 125                                                                                                                         |
| Verify SAP Installation                                        | "How to Verify SAP and the Database Installation with Central Instance" on page 125<br>"How to Verify an SAP Failover Application Server" on page 126                     |
| or                                                             |                                                                                                                                                                           |
| Install SAP, SAP scalable application server, and the database | "How to Install SAP and the Database" on page 120<br>"How to Install an SAP Scalable Application Server" on page 120<br>"Configuring Sun Cluster HA for DBMS" on page 125 |
| Configure the Sun Cluster HA for DBMS                          | "How to Verify an SAP Scalable Application Server" on page 128                                                                                                            |
| Verify SAP Installation                                        |                                                                                                                                                                           |
| Install Sun Cluster HA for SAP packages                        | "How to Install the Sun Cluster HA for SAP Packages" on page 128                                                                                                          |

**TABLE E-2** Task Map: Installing and Configuring Sun Cluster HA for SAP (Continued)

| Task                                                                     | For Instructions, Go To                                                                                                                                                                                                                                                                                                                          |
|--------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Register and configure Sun Cluster HA for SAP as a failover data service | "How to Register and Configure Sun Cluster HA for SAP with Central Instance" on page 136<br>"How to Register and Configure Sun Cluster HA for SAP as a Failover Data Service" on page 137                                                                                                                                                        |
| or                                                                       |                                                                                                                                                                                                                                                                                                                                                  |
| Register and configure Sun Cluster HA for SAP as a scalable data service | "How to Register and Configure Sun Cluster HA for SAP with Central Instance" on page 136<br>"How to Register and Configure Sun Cluster HA for SAP as a Scalable Data Service" on page 138                                                                                                                                                        |
| Set up a lock file                                                       | "Setting Up a Lock File" on page 139                                                                                                                                                                                                                                                                                                             |
| Verify Sun Cluster HA for SAP installation and configuration             | "How to Verify Sun Cluster HA for SAP Installation and Configuration and Central Instance" on page 142<br>"How to Verify the Installation and Configuration of Sun Cluster HA for SAP as a Failover Data Service" on page 143<br>"How to Verify Sun Cluster HA for SAP Installation and Configuration of as a Scalable Data Service" on page 143 |
| Understand Sun Cluster HA for SAP fault monitor                          | "Understanding Sun Cluster HA for SAP Fault Monitor" on page 144                                                                                                                                                                                                                                                                                 |

---

## Planning the Sun Cluster HA for SAP Installation and Configuration

This section contains the information you need to plan your Sun Cluster HA for SAP installation and configuration.

### Configuration Restrictions



---

**Caution** – Your data service configuration might not be supported if you do not observe these restrictions.

---



Use the restrictions in this section to plan the installation and configuration of Sun Cluster HA for SAP. This section provides a list of software and hardware configuration restrictions that apply to Sun Cluster HA for SAP.

For restrictions that apply to all data services, see the *Sun Cluster release notes documentation*.

- **Limit node names as outlined in the SAP installation guide** – This limitation is an SAP software restriction.

## Configuration Requirements



---

**Caution** – Your data service configuration might not be supported if you do not adhere to these requirements.

---

Use the requirements in this section to plan the installation and configuration of Sun Cluster HA for SAP. These requirements apply to Sun Cluster HA for SAP only. You must meet these requirements before you proceed with your Sun Cluster HA for SAP installation and configuration.


For requirements that apply to all data services, see “Configuring and Administering Sun Cluster Data Services” in *Sun Cluster 3.1 Data Service 4/03 Planning and Administration Guide*.

- **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 *Sun Cluster 3.1 Data Service 4/03 for Network File System (NFS)* for details on how to configure Sun Cluster HA for NFS.
- **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.

- **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.


## Standard Data Service Configurations

Use the standard configurations in this section to plan the installation and configuration of Sun Cluster HA for SAP. Sun Cluster HA for SAP supports the standard configurations in this section. Sun Cluster HA for SAP might support additional configurations. However, you must contact your Enterprise Services representative for information on additional configurations.



CLUSTER 1

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




CLUSTER 1

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

---

**Note** – The configuration in Figure E-2 was a common configuration under previous Sun Cluster releases. To use the Sun Cluster software to the fullest extent, configure SAP as shown in Figure E-1 or Figure E-3.

---



CLUSTER 1

**FIGURE E-3** Two-Node Cluster With Central Instance and Development Node

## Configuration Considerations

Use the information in this section to plan the installation and configuration of Sun Cluster HA for SAP. The information in this section encourages you to think about the impact your decisions have on the installation and configuration of Sun Cluster HA for SAP.

### Failover and Scalable Applications

- **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 139.
- **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 consideration 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.

## Scalable Applications

- **Ensure that the `SAPSIDadm` home directory resides on a cluster file system** - This consideration 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.
- **Install the application server's directory locally on each node instead of on a cluster file system** - This consideration ensures that another application server does not overwrite the `log/data/work/sec` directory for the application server.
- **Use the same instance number when you create all application server instances on multiple nodes** - This consideration 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.
- **Place the application servers into multiple resource groups if you want to use the `RGOffload` resource type to shut down one or more application servers when a higher priority resource is failing over** - This consideration provides flexibility and availability if you want to use the `RGOffload` resource type to offload one or more application servers for the database. The value you gain from this consideration 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 *Sun Cluster 3.1 Data Service 4/03 Planning and Administration Guide* for more information on using the `RGOffload` resource type.
- **Create separate scalable application server instances for each SAP logon group.**
- **Create an SAP lock file on the local instance directory** - This consideration prevents a system administrator from manually starting an application instance that is already running.

## Configuration Planning Questions

Use the questions in this section to plan the installation and configuration of Sun Cluster HA for SAP. Insert the answers to these questions into the data service worksheets in the *Sun Cluster release notes documentation*. See "Configuration Considerations" on page 115 for information that might apply to these questions.

- What resource groups will you use for network addresses and application resources and the dependencies between them?
- What is the logical hostname (for failover services) for clients that will access the data service?
- Where will the system configuration files reside?  
See "Determining the Location of the Application Binaries" on page 3 of the *Sun Cluster 3.x Release Notes Supplement* for the advantages and disadvantages of placing the liveCache binaries on the local file system as opposed to the cluster file system.

## Packages and Support

Table E-3 and Table E-4 lists the packages that Sun Cluster HA for SAP supports.

**TABLE E-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. |

The \*\_v2 resource types are the latest version of the resource types (RT) for Sun Cluster HA for SAP. The \*\_v2 resource types are a superset of the original RTs. Whenever possible, use the latest RTs provided.

**TABLE E-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 E-3.                                                                  |
| SUNW.sap_as    | Same as Sun Cluster 3.0 7/01. See Table E-3.                                                                  |
| SUNW.sap_ci_v2 | Added the <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 <code>Network_resources_used</code> resource property to RTR file.                                  |
|                | Added support for scalable application servers.                                                               |
|                | Retained support for failover application servers                                                             |

---

## Upgrading Sun Cluster HA for SAP

As Table E-3 and Table E-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 118 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 118 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

Use this procedure to upgrade a resource type or to convert a failover application server resource to a scalable application server resource. 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. **Register the new resource type.**
5. **Which task are you performing?**
  - If you are upgrading the resource type for the central instance, skip to Step 7.
  - If you are converting a failover application server resource to a scalable application server resource, proceed to Step 6.
6. **Create the new application server resource group.**
7. **Add the scalable application resource to the resource group.**

### Where to Go From Here

Go to “How to Prepare the Nodes” on page 119.

---

## Preparing the Nodes and Disks

This section contains the procedures you need to prepare the nodes and disks.

### ▼ How to Prepare the Nodes

Use this procedure to prepare for the installation and configuration of SAP.

1. **Become superuser on all of the nodes.**
2. **Configure the `/etc/nsswitch.conf` so that Sun Cluster HA for SAP starts and stops correctly in the event of a switchover or a failover.**

On each node that can master the logical host that runs Sun Cluster HA for SAP, include one of the following entries for `group` in the `/etc/nsswitch.conf` file.

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

Sun Cluster HA for SAP uses the `su user` command to start and probe SAP. The network information name service might become unavailable when a cluster node's public network fails. When you add one of the entries for `group` in the `/etc/nsswitch.conf` file, you ensure that the `su(1M)` command does not refer to the NIS/NIS+ name services if the network information name service is unavailable.

### Where to Go From Here

Go to "How to Register and Configure Sun Cluster HA for SAP as a Scalable Data Service" on page 138.

---

## Installing and Configuring SAP and Database

This section contains the procedures you need to install and configure SAP and the database.

## ▼ How to Install SAP and the Database

Use this procedure to install SAP and the database.

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 an SAP Scalable Application Server” on page 120 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 Failover SAP Instances to Run in a Cluster” on page 123 or “How to Install an SAP Scalable Application Server” on page 120.

## ▼ How to Install an SAP Scalable Application Server

Use this procedure to install scalable application server instances. This procedure assumes that you installed the central instance and the database. This procedure includes additional steps for SAP 6.10 and SAP 6.20 users to ensure that Sun Cluster HA for SAP can manage and bring online SAP 6.10 and SAP 6.20 services. SAP 6.10 and SAP 6.20 create one `startsap` script and one `stopsap` script. Other SAP versions create one of each of these scripts for each service you create. This difference accounts for the additional steps for SAP 6.10 and SAP 6.20 users.



---

**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 or the Failover Application Server” on page 140.

- 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.

Example:

```
mkdir -p /usr/sap/local/SID/D03
```



---

**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.

---

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` log file and the `stop` log file will be installed on the local file system.**

Example:

```
ln -s /usr/sap/local/SID/D03 /usr/sap/SID/D03
```

5. Install the application server.
6. Are you using SAP 6.10 or SAP 6.20?
  - If no, skip to Step 11
  - If yes, proceed to Step 7.
7. Become user *sapsidadm*.
8. Make a copy of the **startsap** script and the **stopsap** script, and save these files in the SAPSIDADm home directory. The filenames that you choose specify this instance.

```
cp /usr/sap/SID/SYS/exe/run/startsap \
 $SAPSID_HOME/startsap_instance-number
```

```
cp /usr/sap/SID/SYS/exe/run/stopsap \
 $SAPSID_HOME/stopsap_instance-number
```

9. Make backup copies of the following files because you will modify them. In the SAP profile directory, modify all the filenames for this instance. The filenames that you choose must be specific to this instance, and they must follow the same naming convention you chose in Step 8.

```
mv SAPSID_Service-StringSystem-Number_physical-hostname \
 SAPSID_Service-StringSystem_instance-number
```

```
mv START_Service-StringSystem-Number_physical-hostname \
 START_Service-StringSystem_instance-number
```

10. Modify the contents of the files you created in Step 9 to replace any reference to the physical host with the instance number.



---

**Caution** – It is important that you make your updates consistent so that you can start and stop this application server instance from all the nodes that will run this scalable application server. For example, if you make these changes for SAP instance number 02, then use 02 where this instance number appears. If you do not use a consistent naming convention you will be unable start and stop this application server instance from all the nodes that will run this scalable application server.

---

11. Edit the **start** script and the **stop** script so that the **startup log file** and the **stop log file** will be node specific under the home directories of users *sapsidadm* and *orasapsid*.

Example:

```
vi startsap_D03
```

Before:

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

After:

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

**12. 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`.

**13. Ensure that you can startup and stop the application server from each node, and verify that the log files are in the correct location.**

**14. Create the SAP logon group if you use a logon group.**

## Where to Go From Here

Go to “Configuring Sun Cluster HA for DBMS” on page 125.

## ▼ How to Enable Failover SAP Instances to Run in a Cluster

During SAP installation, the SAP software creates files and shell scripts on the server on which you installed the SAP 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). Use this procedure to enable SAP to run in a cluster.

1. **Make backup copies of the files you will modify in Step 5 through Step 8.**
2. **Log in to the node on which you installed the SAP software.**
3. **Shut down the SAP instances (central instance and application server instances) and the database.**
4. **Are you using SAP 6.10 or SAP 6.20?**
  - If no, skip to Step 6.
  - If yes, proceed to Step 5.
5. **Make a copy of the `startsap` script and the `stopsap` script, and save these files in the `SAPSIDadm` home directory. The filenames that you choose must specify this instance.**

```
cp /usr/sap/SID/SYS/exe/run/startsap \
$SAPSID_HOME/startsap_logical-hostname_instance-number
```

```
cp /usr/sap/SID/SYS/exe/run/startsap \
$SAPSID_HOME/stopsap_logical-hostname_instance-number
```

**6. Become user *sapsidadm*, and then perform the following tasks.**

- In the SAPSIDadm home directory, modify all of the file names that reference a physical server name.
- In the SAPSIDadm 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.

**7. As user *sapsidadm*, add entries for the parameter **SAPLOCALHOST**.**

Add this entry to the *SAPSID\_Service-StringSystem-Number\_logical-hostname* profile file under the */sapmnt/SAPSID/profile* directory.

For Central Instance:

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

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

For Application Server:

```
SAPLOCALHOST=as-logical-hostname
```

**8. 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.

**9. Ensure that the */usr/sap/tmp* directory owned by user *sapsidadm* and group *sapsys* exists on all nodes that can master the failover SAP instance.**

## Where to Go From Here

Go to “Configuring Sun Cluster HA for DBMS” on page 125.

---

## 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 *Sun Cluster 3.1 Data Service 4/03 for Oracle* 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.HAStoragePlus` resource type if you use Oracle. See the procedure “Synchronizing the Startups Between Resource Groups and Disk Device Groups” in *Sun Cluster 3.1 Data Service 4/03 Planning and Administration Guide* for more information.

### Where to Go From Here

Go to “How to Verify SAP and the Database Installation with Central Instance” on page 125 or “How to Verify an SAP Scalable Application Server” on page 128.

---

## Verifying the SAP Installation

This section contains the procedures you need to verify the SAP installation.

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

Use this procedure to verify SAP central instance. Perform the following steps on all of the potential nodes on which the central instance can run.

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

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

---

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

---

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. **Ensure that the database is running.**
7. **Manually start the central instance.**
8. **Start the SAP GUI using the logical hostname, and verify that SAP initializes correctly.**

The default dispatcher port is 3200.
9. **Manually stop the central instance.**
10. **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 an SAP Failover Application Server” on page 126.

## ▼ How to Verify an SAP Failover Application Server

Use this procedure to verify SAP and the database installation for the failover application server. Perform the following steps 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-fo-resource-group
```

5. Log in to the cluster member that hosts the application server resource group.

6. Manually start the application server.

7. Start the SAP GUI using the logical hostname, and verify that SAP initializes correctly.

8. Manually stop the application server.

9. 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 “How to Install the Sun Cluster HA for SAP Packages” on page 128.

## How to Verify an SAP Scalable Application Server

If you installed scalable application server instances in “How to Install an SAP Scalable Application Server” on page 120, you verified the installation of an SAP scalable application server in Step 13 of “How to Install an SAP Scalable Application Server” on page 120.

### Where to Go From Here

Go to “How to Install the Sun Cluster HA for SAP Packages” on page 128.

---

## Installing the Sun Cluster HA for SAP Packages

This section contains the procedure you need to install the Sun Cluster HA for SAP packages.

### ▼ How to Install the Sun Cluster HA for SAP Packages

Use this procedure to install the Sun Cluster HA for SAP packages. You need the Sun Cluster 3.1 Agents CD-ROM to perform this procedure. This procedure assumes that you did not install the data service packages during your initial Sun Cluster installation.

**1. Load the Sun Cluster 3.1 Agents CD-ROM into the CD-ROM drive.**

**2. Run the `scinstall` utility with no options.**

This step starts the `scinstall` utility in interactive mode.

**3. Choose the Add Support for New Data Service to This Cluster Node menu option.**

The `scinstall` utility prompts you for additional information.

**4. Provide the path to the Sun Cluster 3.1 Agents CD-ROM.**

The utility refers to the CD-ROM as the “data services cd.”

**5. Specify the data service to install.**

The `scinstall` utility lists the data service that you selected and asks you to confirm your choice.



6. Exit the `scinstall` utility.
7. Unload the CD-ROM from the drive.

## Where to Go From Here

Go to “Registering and Configuring Sun Cluster HA for SAP” on page 129.

---

# Registering and Configuring Sun Cluster HA for SAP

This section contains the procedures you need to configure Sun Cluster HA for SAP.

## Sun Cluster HA for SAP Extension Properties

Use the extension properties in Table E-5 and Table E-6 to create your resources. Use the command line `scrgadm -x parameter=value` to configure extension properties when you create your resource. Use the procedure in “Administering Data Service Resources” in *Sun Cluster 3.1 Data Service 4/03 Planning and Administration Guide* to configure the extension properties if you have already created your resources. You can update some extension properties dynamically. You can update others, however, only when you create or disable a resource. The Tunable entries indicate when you can update each property. See Appendix A for details on all Sun Cluster properties.

**TABLE E-5** Sun Cluster HA for SAP Extension Properties for the Central Instance

| Property Category | Property Name           | Description                                                                                                                                                                |
|-------------------|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SAP Configuration | SAPSID                  | SAP system ID or SID.<br><b>Default:</b> None<br><b>Tunable:</b> When disabled                                                                                             |
|                   | Ci_instance_id          | Two-digit SAP system number.<br><b>Default:</b> 00<br><b>Tunable:</b> When disabled                                                                                        |
|                   | Ci_services_string      | String of central instance services.<br><b>Default:</b> DVEBMGS<br><b>Tunable:</b> When disabled                                                                           |
| Starting SAP      | Ci_start_retry_interval | The interval in seconds to wait between attempting to connect to the database before starting the central instance.<br><b>Default:</b> 30<br><b>Tunable:</b> When disabled |
|                   | Ci_startup_script       | Name of the SAP startup script for this instance in your <i>SID</i> adm home directory.<br><b>Default:</b> None<br><b>Tunable:</b> When disabled                           |

**TABLE E-5** Sun Cluster HA for SAP Extension Properties for the Central Instance  
(Continued)

| Property Category | Property Name      | Description                                                                                                                                                                                                                                                                               |
|-------------------|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Stopping SAP      | stop_sap_pct       | <p>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.</p> <p><b>Default:</b> 95</p> <p><b>Tunable:</b> When disabled</p> |
|                   | ci_shutdown_script | <p>Name of the SAP shutdown script for this instance in your <i>SIDadm</i> home directory.</p> <p><b>Default:</b> None</p> <p><b>Tunable:</b> When disabled</p>                                                                                                                           |

**TABLE E-5** Sun Cluster HA for SAP Extension Properties for the Central Instance  
(Continued)

| Property Category | Property Name                 | Description                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Probe             | Message_server_name           | The name of the SAP Message Server.<br><br><b>Default:</b> sapms<br>SAPSID<br><br><b>Tunable:</b> When disabled                                                                                                                                                                                                                                                                                                        |
|                   | Lgtst_ms_with_logicalhostname | How to check the SAP Message Server with the SAP lgtst utility. The lgtst 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 local host (loopback) name. If you set this resource property to TRUE, use a logical hostname. Otherwise, use a localhost name.<br><br><b>Default:</b> TRUE<br><b>Tunable:</b> 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.<br><br><b>Default:</b> 2<br><b>Tunable:</b> When disabled                                                                                                                                                                                                                   |
|                   | Probe_timeout                 | Timeout value in seconds for the probes.<br><br><b>Default:</b> 120<br><b>Tunable:</b> Any time                                                                                                                                                                                                                                                                                                                        |
|                   | Monitor_retry_count           | Number of PMF restarts that are allowed for the fault monitor.<br><br><b>Default:</b> 4<br><b>Tunable:</b> Any time                                                                                                                                                                                                                                                                                                    |
|                   | Monitor_retry_interval        | Time interval in minutes for the fault monitor restarts.<br><br><b>Default:</b> 2<br><b>Tunable:</b> Any time                                                                                                                                                                                                                                                                                                          |

**TABLE E-5** Sun Cluster HA for SAP Extension Properties for the Central Instance  
(Continued)

| Property Category  | Property Name       | Description                                                                                                                                                                                           |
|--------------------|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Development System | Shutdown_dev        | Whether the RGM should shut down the development system before starting up the central instance.<br><br><b>Default:</b> FALSE<br><b>Tunable:</b> 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).<br><br><b>Default:</b> None<br><b>Tunable:</b> 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.<br><br><b>Default:</b> None<br><b>Tunable:</b> 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.<br><br><b>Default:</b> 20<br><b>Tunable:</b> When disabled       |

**TABLE E-6** Sun Cluster HA for SAP Extension Properties for the Application Servers

| Property Category | Property Name        | Description                                                                                                                                                               |
|-------------------|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SAP Configuration | SAPSID               | SAP system name or <i>SAPSID</i> for the application server.<br><b>Default:</b> None<br><b>Tunable:</b> When disabled                                                     |
|                   | As_instance_id       | Two-digit SAP system number for the application server.<br><b>Default:</b> None<br><b>Tunable:</b> When disabled                                                          |
|                   | As_services_string   | String of application server services.<br><b>Default:</b> D<br><b>Tunable:</b> When disabled                                                                              |
| Starting SAP      | As_db_retry_interval | The interval in seconds to wait between attempting to connect to the database and starting the application server.<br><b>Default:</b> 30<br><b>Tunable:</b> When disabled |
|                   | As_startup_script    | Name of the SAP startup script for the application server.<br><b>Default:</b> None<br><b>Tunable:</b> When disabled                                                       |

**TABLE E-6** Sun Cluster HA for SAP Extension Properties for the Application Servers  
(Continued)

| Property Category | Property Name          | Description                                                                                                                                                                                                                                                                    |
|-------------------|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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.<br><br><b>Default:</b> 95<br><b>Tunable:</b> When disabled |
|                   | As_shutdown_script     | Name of the SAP shutdown script for the application server.<br><br><b>Default:</b> None<br><b>Tunable:</b> When disabled                                                                                                                                                       |
| Probe             | Probe_timeout          | Time-out value in seconds for the probes.<br><br><b>Default:</b> 60<br><b>Tunable:</b> Any time                                                                                                                                                                                |
|                   | Monitor_retry_count    | Number of PMF restarts that the probe allows for the fault monitor.<br><br><b>Default:</b> 4<br><b>Tunable:</b> Any time                                                                                                                                                       |
|                   | Monitor_retry_interval | Time interval in minutes for fault monitor restarts.<br><br><b>Default:</b> 2<br><b>Tunable:</b> Any time                                                                                                                                                                      |

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

Use this procedure to configure Sun Cluster HA for SAP with central instance.

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. **Add the HAStoragePlus resource to the central instance resource group.**

```
scrgadm -a -t SUNW.HAStoragePlus
scrgadm -a -j ci-storage-resource \
-g sap-ci-resource-group \
-t SUNW.HAStoragePlus -x filesystemmountpoints=mountpoint, ... |
```

For more details on how to set up an HAStoragePlus resource, see “Enabling Highly Available Local File Systems” in *Sun Cluster 3.1 Data Service 4/03 Planning and Administration Guide*.

4. **Enable the central instance storage resource.**

```
scswitch -e -j ci-storage-resource
```

5. **Create SAP central instance 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_instance_id=ci-instance-id \
-x Ci_startup_script=ci-startup-script \
-x Ci_shutdown_script=ci-shutdown-script \
-y resource_dependencies=ci-storage-resource
```

See *Sun Cluster 3.1 Data Service 4/03 Planning and Administration Guide* for a list of extension properties.

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

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

If you configure the central instance resource to shut down a development system, you will receive the following console message.

```
ERROR : SAPSYSTEMNAME not set
Please check environment and restart
```

This message displays when the central instance starts on a node that does not have the development system installed and that is not meant to run the central instance. SAP renders this message, and you can safely ignore it.



## Where to Go From Here

Go to “How to Register and Configure Sun Cluster HA for SAP as a Failover Data Service” on page 137 or “How to Register and Configure Sun Cluster HA for SAP as a Scalable Data Service” on page 138.

### ▼ How to Register and Configure Sun Cluster HA for SAP as a Failover Data Service

Use this procedure to configure Sun Cluster HA for SAP as a failover data service.

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. **Add the HAStoragePlus resource to the failover application server resource group.**

```
scrgadm -a -t SUNW.HAStoragePlus
scrgadm -a -j sap-as-storage-resource -g sap-as-fo-resource-group \
-t SUNW.HAStoragePlus \
-x filesystemmountpoints=mountpoint, ...
```

For more details on how to set up an HAStoragePlus resource, see “Enabling Highly Available Local File Systems” in *Sun Cluster 3.1 Data Service 4/03 Planning and Administration Guide*.

4. **Enable the failover application server storage resource.**

```
scswitch -e -j sap-as-storage-resource
```

5. **Create SAP application server resources in their 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 \
-y resource_dependencies=sap-as-storage-resource
```

See “Sun Cluster HA for SAP Extension Properties” on page 129 for a list of extension properties.

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

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



5. Enable the failover application server storage resource.

```
scswitch -e -j sap-as-storage-resource
```

6. 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 \
-y resource_dependencies=sap-as-storage-resource
```

See “Sun Cluster HA for SAP Extension Properties” on page 129 for a list of extension properties.

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

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

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

- If you use the 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 Sun Cluster HA for SAP Installation and Configuration and Central Instance” on page 142.

---

## 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 or the failover application server.
- 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 `/var/adm/messages`.

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 or the Failover Application Server

Use this procedure to set up a lock file for the central instance or the failover application server.

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 or the failover application server lock file on a *cluster* file system.**
3. **Edit the profile that `sapstart` uses to start the instance such that you add the new SAP parameter, `sapstart/lockfile`, for central instance or failover application server. This profile is the one that is passed to `sapstart` as a parameter in the `startsap` script.**

For central instance, enter the following.

```
sapstart/lockfile =/usr/sap/SID/ Service-StringSystem-Number/work/startup_lockfile
```

For failover application server, enter the following.

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

```
sapstart/lockfile
New parameter name.
```

```
/usr/sap/SID/Service-StringSystem-Number/work
Work directory for the central instance.
```

```
/usr/sap/SID/Dinstance-id/work
Work directory for failover application server.
```

startup\_lockfile

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, a startup of the same instance from multiple nodes cannot be prevented.

---

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

Use this procedure to set up a lock file for a scalable application server.

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 that `sapstart` uses to start the instance such that you add the new SAP parameter, `sapstart/lockfile`, for scalable application server. This profile is the one that is passed to `sapstart` as a parameter in the `startsap` script.**

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

sapstart/lockfile                      New parameter name.

/usr/sap/local/SID/Dinstance-id/work      Work directory for the scalable application server.

startup\_lockfile                      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 on the *same* node.

---

---

## Verifying the Sun Cluster HA for SAP Installation and Configuration

This section contains the procedure you need to verify that you installed and configured your data service correctly.

### ▼ How to Verify Sun Cluster HA for SAP Installation and Configuration and Central Instance

Use this procedure to verify the Sun Cluster HA for SAP installation and configuration and central instance.

1. **Log in to the node that hosts the resource group that contains the SAP central instance resource.**
2. **Start the SAP GUI to check that Sun Cluster HA for SAP is functioning correctly.**
3. **As user *sapsidadm*, use the central instance *stopsap* script to shut down the SAP central instance.**

The Sun Cluster software restarts the central instance.

4. **As user *root*, switch the SAP resource group to another cluster member.**

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

5. **Verify that the SAP central instance starts on this node.**
6. **Repeat Step 1 through Step 5 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 as a Failover Data Service” on page 143 or “How to Verify Sun Cluster HA for SAP Installation and Configuration of as a Scalable Data Service” on page 143.

## ▼ How to Verify the Installation and Configuration of Sun Cluster HA for SAP as a Failover Data Service

Use this procedure to verify the installation and configuration of Sun Cluster HA for SAP as a failover data service.

1. **Log in to the node that currently hosts the resource group that contains the SAP application server resource.**
2. **As user *sapsidadm*, start the SAP GUI to check that the application server is functioning correctly.**
3. **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 restarts the application server.
4. **As user *root*, switch the resource group that contains the SAP application server resource to another cluster member.**  

```
scswitch -z -h node2 -g sap-as-resource-group
```
5. **Verify that the SAP application server starts on the node you identified in Step 4.**
6. **Repeat Step 1 through Step 5 until you have tested all of the potential nodes on which the SAP application server can run.**

## ▼ How to Verify Sun Cluster HA for SAP Installation and Configuration of as a Scalable Data Service

Use this procedure to verify the installation and configuration of Sun Cluster HA for SAP as a scalable data service.

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 restarts the application server.
5. **Repeat Step 1 through Step 3 until you have tested all of the potential nodes on which the SAP application server can run.**

---

## Understanding Sun Cluster HA for SAP Fault Monitor

The Sun Cluster HA for SAP fault monitor checks SAP process and database availability. SAP process availability 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 availability, SAP database availability uses has no impact on SAP resources' failure history. Database availability does, however, trigger the SAP fault monitor to log any `syslog` messages to `/var/adm/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 availability 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. **Availability 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 local hostname (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 database availability. If the database connection status check fails, the fault monitor logs the message, `Database might be down`, to `/var/adm/messages`. 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 `/var/adm/messages` 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 availability 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. **Availability 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 database availability. If the database connection status check fails, the fault monitor logs the message, `Database might be down`, to `/var/adm/messages` 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 `/var/adm/messages`. The fault monitor then sets the status of the SAP resource to `OK`.

#### 4. Evaluate 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.

---

## Upgrading Sun Cluster Software

---

This appendix provides the following step-by-step procedures to upgrade a Sun Cluster 3.0 configuration to Sun Cluster 3.1 software, including upgrade from Solaris 8 to Solaris 9 software, or to upgrade a Sun Cluster 3.1 configuration that runs on Solaris 8 software to Solaris 9 software:

- “How to Prepare the Cluster for Upgrade” on page 149
- “How to Upgrade the Solaris Operating Environment” on page 152
- “How to Upgrade to Sun Cluster 3.1 Software” on page 154
- “How to Upgrade Sun Cluster-Module Software for Sun Management Center” on page 160
- “How to Finish Upgrading to Sun Cluster 3.1 Software” on page 161
- “How to Handle Storage Reconfiguration During an Upgrade” on page 162
- “How to Resolve Mistaken Storage Changes During an Upgrade” on page 163

This appendix replaces the section “Upgrading From Sun Cluster 3.0 to Sun Cluster 3.1 Software” in *Sun Cluster 3.1 Software Installation Guide*.

---

## Upgrading to Sun Cluster 3.1 Software

Perform the following tasks to upgrade from Sun Cluster 3.0 software to Sun Cluster 3.1 software, including upgrade from Solaris 8 to Solaris 9 software, or to upgrade a Sun Cluster 3.1 configuration that runs on Solaris 8 software to Solaris 9 software.

**TABLE F-1** Task Map: Upgrading to Sun Cluster 3.1 Software

| Task                                               | Instructions                                        |
|----------------------------------------------------|-----------------------------------------------------|
| 1. Read the upgrade requirements and restrictions. | “Upgrade Requirements and Restrictions” on page 148 |

**TABLE F-1** Task Map: Upgrading to Sun Cluster 3.1 Software (Continued)

| Task                                                                                                                                                                                                                          | Instructions                                                                       |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------|
| 2. Take the cluster out of production, disable resources, and back up shared data and system disks.                                                                                                                           | "How to Prepare the Cluster for Upgrade" on page 149                               |
| 3. Upgrade the Solaris software, if necessary, to a supported Solaris update release. Optionally, upgrade VERITAS Volume Manager (VxVM).                                                                                      | "How to Upgrade the Solaris Operating Environment" on page 152                     |
| 4. Upgrade to Sun Cluster 3.1 framework and data-service software. This is required for upgrade from Solaris 8 software to Solaris 9 software. If necessary, upgrade applications. If you upgraded VxVM, upgrade disk groups. | "How to Upgrade to Sun Cluster 3.1 Software" on page 154                           |
| 5. (Optional) Upgrade the Sun Cluster module to Sun Management Center, if needed.                                                                                                                                             | "How to Upgrade Sun Cluster-Module Software for Sun Management Center" on page 160 |
| 6. Reregister resource types, enable resources, and bring resource groups online.                                                                                                                                             | "How to Finish Upgrading to Sun Cluster 3.1 Software" on page 161                  |

## Upgrade Requirements and Restrictions

Observe the following requirements and restrictions when you upgrade to Sun Cluster 3.1 software:

- The cluster must run on or be upgraded to at least Solaris 8 2/02 software, including the most current required patches.
- The cluster hardware must be a supported configuration for Sun Cluster 3.1 software. Contact your Sun representative for information about current supported Sun Cluster configurations.
- You must upgrade all software to a version that is supported by Sun Cluster 3.1 software. For example, you must upgrade a data service that is supported on Sun Cluster 3.0 software but is not supported on Sun Cluster 3.1 software to the version of that data service that is supported on Sun Cluster 3.1 software. If the related application is not supported on Sun Cluster 3.1 software, you must also upgrade that application to a supported release.
- The `scinstall` upgrade utility only upgrades those data services that are provided with Sun Cluster 3.1 software. You must manually upgrade any custom or third-party data services.
- Have available the test IP addresses to use with your public network adapters when NAFO groups are converted to Internet Protocol (IP) Network Multipathing groups. The `scinstall` upgrade utility prompts you for a test IP address for each public network adapter in the cluster. A test IP address must be on the same subnet as the primary IP address for the adapter.

See the *IP Network Multipathing Administration Guide* (Solaris 8) or *System Administration Guide: IP Services* (Solaris 9) for information about test IP addresses for IP Network Multipathing groups.

- Sun Cluster 3.1 software supports direct upgrade only from Sun Cluster 3.x software.
- Sun Cluster 3.1 software does not support any downgrade of Sun Cluster software.

## ▼ How to Prepare the Cluster for Upgrade

Before you upgrade the software, perform the following steps to take the cluster out of production:

1. **Ensure that the configuration meets requirements for upgrade.**  
See “Upgrade Requirements and Restrictions” on page 148.
2. **Have available the CD-ROMs, documentation, and patches for all software products you are upgrading.**
  - Solaris 8 or Solaris 9 operating environment
  - Sun Cluster 3.1 framework
  - Sun Cluster 3.1 data services (agents)
  - Applications that are managed by Sun Cluster 3.1 data-service agents
  - VERITAS Volume Manager
  - Patch 113801-01 or later, which is required to upgrade from Solaris 8 software to Solaris 9 software

See “Patches and Required Firmware Levels” in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.

3. **(Optional) Install Sun Cluster 3.1 documentation.**  
Install the documentation packages on your preferred location, such as an administrative console or a documentation server. See the `index.html` file at the top level of the Sun Cluster 3.1 CD-ROM to access installation instructions.
4. **Have available your list of test IP addresses, one for each public network adapter in the cluster.**  
A test IP address is required for each public network adapter in the cluster, regardless of whether the adapter is the active adapter or the backup adapter in a NAFO group. The test IP addresses will be used to reconfigure the adapters to use IP Network Multipathing.

---

**Note** – Each test IP address must be on the same subnet as the existing IP address that is used by the public network adapter.

---

To list the public network adapters on a node, run the following command:

```
% pnmstat
```

See the *IP Network Multipathing Administration Guide* (Solaris 8) or *System Administration Guide: IP Services* (Solaris 9) for more information about test IP addresses for IP Network Multipathing.

**5. Notify users that cluster services will be unavailable during upgrade.**

**6. Ensure that the cluster is functioning normally.**

- To view the current status of the cluster, run the following command from any node:

```
% scstat
```

See the `scstat(1M)` man page for more information.

- Search the `/var/adm/messages` log on the same node for unresolved error messages or warning messages.
- Check volume manager status.

**7. Become superuser on a node of the cluster.**

**8. Switch each resource group offline.**

```
scswitch -F -g resource-group
```

`-F` Switches a resource group offline

`-g resource-group` Specifies the name of the resource group to take offline

**9. Disable all resources in the cluster.**

The disabling of resources before upgrade prevents the cluster from bringing the resources online automatically if a node is mistakenly rebooted into cluster mode.

---

**Note** – If you are upgrading from a Sun Cluster 3.1 release, you can use the `scsetup(1M)` utility instead of the command line. From the Main Menu, choose Resource Groups, then choose Enable/Disable Resources.

---

**a. From any node, list all enabled resources in the cluster.**

```
scrgadm -pv | grep "Res enabled"
```

**b. Identify those resources that depend on other resources.**

You must disable dependent resources first before you disable the resources that they depend on.

**c. Disable each enabled resource in the cluster.**

```
scswitch -n -j resource
```

-n Disables  
-j *resource* Specifies the resource  
See the `scswitch(1M)` man page for more information.

**10. Move each resource group to the unmanaged state.**

```
scswitch -u -g resource-group
```

-u Moves the specified resource group to the unmanaged state

-g *resource-group* Specifies the name of the resource group to move into the unmanaged state

**11. Verify that all resources on all nodes are disabled and that all resource groups are in the unmanaged state.**

```
scstat -g
```

**12. Stop all databases that are running on each node of the cluster.**

**13. Ensure that all shared data is backed up.**

**14. From one node, shut down the cluster.**

```
scshutdown
ok
```

See the `scshutdown(1M)` man page for more information.

**15. Boot each node into noncluster mode.**

```
ok boot -x
```

**16. Ensure that each system disk is backed up.**

**17. Determine whether to upgrade the Solaris operating environment.**

- If Sun Cluster 3.1 software does not support the release of the Solaris environment that you currently run on your cluster, you must upgrade the Solaris software to a supported release. Go to “How to Upgrade the Solaris Operating Environment” on page 152.
- If your cluster configuration already runs on a release of the Solaris environment that supports Sun Cluster 3.1 software, go to “How to Upgrade to Sun Cluster 3.1 Software” on page 154.

See “Supported Products” in *Sun Cluster 3.1 Release Notes* for more information.

## ▼ How to Upgrade the Solaris Operating Environment

Perform this procedure on each node in the cluster to upgrade the Solaris operating environment.

---

**Note** – The cluster must already run on, or be upgraded to, at least the minimum required level of the Solaris 8 or Solaris 9 environment to support Sun Cluster 3.1 software. See “Supported Products” in *Sun Cluster 3.1 Release Notes* for more information.

---

1. Ensure that all steps in “How to Prepare the Cluster for Upgrade” on page 149 are completed.
2. Become superuser on the cluster node to upgrade.
3. Determine whether the following Apache links already exist, and if so, whether the file names contain an uppercase K or S:

```
/etc/rc0.d/K16apache
/etc/rc1.d/K16apache
/etc/rc2.d/K16apache
/etc/rc3.d/S50apache
/etc/rcS.d/K16apache
```

- If these links already exist and do contain an uppercase K or S in the file name, no further action is necessary for these links.
  - If these links do not exist, or if these links exist but instead contain a lowercase k or s in the file name, you move aside these links in Step 8.
4. Comment out all entries for globally mounted file systems in the `/etc/vfstab` file.
    - a. Make a record of all entries that are already commented out for later reference.
    - b. Temporarily comment out all entries for globally mounted file systems in the `/etc/vfstab` file.

Entries for globally mounted file systems contain the `global` mount option. Comment out these entries to prevent the Solaris upgrade from attempting to mount the global devices.
  5. Determine which procedure to follow to upgrade the Solaris operating environment.



| Volume Manager                             | Procedure to Use                    | Instructions                                      |
|--------------------------------------------|-------------------------------------|---------------------------------------------------|
| Solstice DiskSuite™/Solaris Volume Manager | Upgrading Solaris software          | Solaris 8 or Solaris 9 installation documentation |
| VERITAS Volume Manager                     | Upgrading VxVM and Solaris software | VERITAS Volume Manager installation documentation |

**6. Upgrade the Solaris software, following the procedure you selected in Step 5.**

---

**Note** – Ignore the instruction to reboot at the end of the Solaris software upgrade process. You must first perform Step 7 and Step 8, then reboot into noncluster mode in Step 9 to complete Solaris software upgrade.

If you are instructed to reboot a node at other times in the upgrade process, always add the `-x` option to the command. This option ensures that the node reboots into noncluster mode. For example, either of the following two commands boot a node into single-user noncluster mode:

```
reboot -- -xs
ok boot -xs
```

---

**7. In the `/a/etc/vfstab` file, uncomment those entries for globally mounted file systems that you commented out in Step 4.**

**8. If the Apache links in Step 3 did not already exist or if they contained a lowercase `k` or `s` in the file names *before* you upgraded the Solaris software, move aside the restored Apache links.**

Use the following commands to rename the files with a lowercase `k` or `s`:

```
mv /a/etc/rc0.d/K16apache /a/etc/rc0.d/k16apache
mv /a/etc/rc1.d/K16apache /a/etc/rc1.d/k16apache
mv /a/etc/rc2.d/K16apache /a/etc/rc2.d/k16apache
mv /a/etc/rc3.d/S50apache /a/etc/rc3.d/s50apache
mv /a/etc/rcS.d/K16apache /a/etc/rcS.d/k16apache
```

**9. Reboot the node into noncluster mode.**

Include the double dashes (`--`) in the following command:

```
reboot -- -x
```

**10. Install any required Solaris software patches and hardware-related patches, and download any needed firmware that is contained in the hardware patches.**

For Solstice DiskSuite software (Solaris 8), also install any Solstice DiskSuite software patches.

---

**Note** – Do not reboot after you add patches. You reboot the node after you upgrade the Sun Cluster software.

---

See “Patches and Required Firmware Levels” in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.

**11. Upgrade to Sun Cluster 3.1 software.**

Go to “How to Upgrade to Sun Cluster 3.1 Software” on page 154.

---

**Note** – To complete upgrade from Solaris 8 to Solaris 9 software, you must also upgrade to the Solaris 9 version of Sun Cluster 3.1 software, even if the cluster already runs on Sun Cluster 3.1 software.

---

## ▼ How to Upgrade to Sun Cluster 3.1 Software

This procedure describes how to upgrade the cluster to Sun Cluster 3.1 software. You must also perform this procedure to complete cluster upgrade from Solaris 8 to Solaris 9 software.

---

**Tip** – You can perform this procedure on more than one node at the same time.

---

**1. Ensure that all steps in “How to Prepare the Cluster for Upgrade” on page 149 are completed.**

If you upgraded from Solaris 8 to Solaris 9 software, also ensure that all steps in “How to Upgrade the Solaris Operating Environment” on page 152 are completed.

**2. Become superuser on a node of the cluster.**

**3. Ensure that you have installed all required Solaris software patches and hardware-related patches.**

For Solstice DiskSuite software (Solaris 8), also ensure that you have installed all required Solstice DiskSuite software patches.

**4. Insert the Sun Cluster 3.1 CD-ROM into the CD-ROM drive on the node.**

If the Volume Management daemon `vold(1M)` is running and configured to manage CD-ROM devices, it automatically mounts the CD-ROM on the `/cdrom/suncluster_3_1` directory.

**5. Upgrade the node to Sun Cluster 3.1 software.**

- a. Change to the `/cdrom/suncluster_3_1/SunCluster_3.1/Sol_ver/Tools` directory, where *ver* is 8 (for Solaris 8) or 9 (for Solaris 9) .

```
cd /cdrom/suncluster_3_1/SunCluster_3.1/Sol_ver/Tools
```

- b. Upgrade the cluster framework software.

- To upgrade from Sun Cluster 3.0 software, run the following command:

```
./scinstall -u update -S interact
```

-S Specifies the test IP addresses to use to convert NAFO groups to IP Network Multipathing groups

interact Specifies that `scinstall` prompts the user for each test IP address needed

- To upgrade from Sun Cluster 3.1 software, run the following command:

```
./scinstall -u update
```

---

**Tip** – If upgrade processing is interrupted, use the `scstat(1M)` command to ensure that the node is in noncluster mode (Offline), then restart the `scinstall` command.

```
scstat -n
-- Cluster Nodes --
 Node name Status
 ----- -
Cluster node: nodename Offline
Cluster node: nodename Offline
```

---

See the `scinstall(1M)` man page for more information. See the *IP Network Multipathing Administration Guide* (Solaris 8) or *System Administration Guide: IP Services* (Solaris 9) for information about test addresses for IP Network Multipathing.

---

**Note** – Sun Cluster 3.1 software requires at least version 3.5.1 of Sun Explorer software. Upgrade to Sun Cluster software includes installing Sun Explorer data collector software, to be used in conjunction with the `sccheck` utility. If another version of Sun Explorer software was already installed before Sun Cluster upgrade, it is replaced by the version that is provided with Sun Cluster software. Options such as user identity and data delivery are preserved, but `crontab` entries must be manually re-created.

---

During Sun Cluster upgrade, `scinstall` might make one or more of the following configuration changes:

- Convert NAFO groups to IP Network Multipathing groups but keep the original NAFO-group name.
- Rename the `ntp.conf` file to `ntp.conf.cluster`, if `ntp.conf.cluster` does not already exist on the node.
- Set the `local-mac-address?` variable to `true`, if the variable is not already set to that value.

c. **Change to the CD-ROM root directory and eject the CD-ROM.**

d. **Install any Sun Cluster 3.1 patches.**

---

**Note** – If you upgraded from Solaris 8 software to Solaris 9 software, install Patch 113801–01 or later before you proceed to the next step.

---

See “Patches and Required Firmware Levels” in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.

---

**Note** – Do not reboot the node at this time.

---

6. **Upgrade software applications that are installed on the cluster and apply application patches as needed.**

Ensure that application levels are compatible with the current version of Sun Cluster and Solaris software. See your application documentation for installation instructions. In addition, follow these guidelines to upgrade applications in a Sun Cluster 3.1 configuration:

- If the applications are stored on shared disks, you must master the relevant disk groups and manually mount the relevant file systems before you upgrade the application.
- If you are instructed to reboot a node during the upgrade process, always add the `-x` option to the command. This option ensures that the node reboots into noncluster mode. For example, either of the following two commands boot a node into single-user noncluster mode:

```
reboot -- -xs
ok boot -xs
```

7. **Upgrade Sun Cluster data services to the Sun Cluster 3.1 software versions.**

---

**Note** – Only those data services that are provided on the Sun Cluster 3.1 Agents CD-ROM are automatically upgraded by `scinstall(1M)`. You must manually upgrade any custom or third-party data services.

---

a. **Insert the Sun Cluster 3.1 Agents CD-ROM into the CD-ROM drive on the node to upgrade.**

b. **Upgrade the data-service software.**

```
scinstall -u update -s all -d /cdrom/cdrom0
```

-u update            Specifies upgrade

-s all                Updates all Sun Cluster data services that are installed on the node

---

**Tip** – If upgrade processing is interrupted, use the `scstat(1M)` command to ensure that the node is in noncluster mode (`Offline`), then restart the `scinstall` command.

```
scstat -n
```

```
-- Cluster Nodes --
```

|               | Node name       | Status  |
|---------------|-----------------|---------|
|               | -----           | -----   |
| Cluster node: | <i>nodename</i> | Offline |
| Cluster node: | <i>nodename</i> | Offline |

---

c. **Change to the CD-ROM root directory and eject the CD-ROM.**

d. **As needed, manually upgrade any custom data services that are not supplied on the Sun Cluster 3.1 Agents CD-ROM.**

e. **Install any Sun Cluster 3.1 data-service patches.**

See “Patches and Required Firmware Levels” in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.

8. **After all nodes are upgraded, reboot each node into the cluster.**

```
reboot
```

9. **Verify that all upgraded software is at the same version on all upgraded nodes.**

a. **On each upgraded node, view the installed levels of Sun Cluster software.**

```
scinstall -pv
```

- b. From one node, verify that all upgraded cluster nodes are running in cluster mode (Online).

```
scstat -n
```

See the `scstat(1M)` man page for more information about displaying cluster status.

**10. Did you upgrade from Solaris 8 to Solaris 9 software?**

- If no, skip to Step 14.
- If yes, proceed to Step 11.

**11. On each node, run the following command to verify the consistency of the storage configuration:**

```
scdidadm -c
```

```
-c Perform a consistency check
```




---

**Caution** – Do not proceed to Step 12 until your configuration passes this consistency check. Failure to do so might result in errors in device identification and cause data corruption.

---

The following table lists the possible output from the `scdidadm -c` command and the action you must take, if any.

| Example Message                                                                                                    | Action to Take                                                                                                                 |
|--------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| device id for 'phys-schost-1:/dev/rdisk/c1t3d0' does not match physical device's id, device may have been replaced | Go to "Recovering From Storage Configuration Changes During Upgrade" on page 162 and perform the appropriate repair procedure. |
| device id for 'phys-schost-1:/dev/rdisk/c0t0d0' needs to be updated, run <code>scdidadm -R</code> to update        | None. You update this device ID in Step 12.                                                                                    |
| No output message                                                                                                  | None.                                                                                                                          |

See the `scdidadm(1M)` man page for more information.

**12. On each node, migrate the Sun Cluster storage database to Solaris 9 device IDs.**

```
scdidadm -R all
```

```
-R Perform repair procedures
```

```
all Specify all devices
```

**13. On each node, run the following command to verify that storage database migration to Solaris 9 device IDs is successful:**

```
scdidadm -c
```

- If the `scdidadm` command displays a message, return to Step 11 to make further corrections to the storage configuration or the storage database.
- If the `scdidadm` command displays no messages, the device-ID migration is successful. If device-ID migration is verified on all cluster nodes, proceed to Step 14.

**14. Did you upgrade VxVM?**

- If no, proceed to Step 15.
- If yes, upgrade all disk groups.

To upgrade a disk group to the highest version supported by the VxVM release you installed, run the following command from the primary node of the disk group:

```
vxdg upgrade dgname
```

See your VxVM administration documentation for more information about upgrading disk groups.

**15. Do you intend to use Sun Management Center to monitor the cluster?**

- If yes, go to “How to Upgrade Sun Cluster-Module Software for Sun Management Center” on page 160.
- If no, go to “How to Finish Upgrading to Sun Cluster 3.1 Software” on page 161.

## Example—Upgrade From Sun Cluster 3.0 to Sun Cluster 3.1 Software

The following example shows the process of upgrading a two-node cluster, including data services, from Sun Cluster 3.0 to Sun Cluster 3.1 software on the Solaris 8 operating environment. The cluster node names are `phys-schost-1` and `phys-schost-2`.

*(On the first node, upgrade framework software from the Sun Cluster 3.1 CD-ROM)*

```
phys-schost-1# cd /cdrom/suncluster_3_1/SunCluster_3.1/Sol_8/Tools
phys-schost-1# ./scinstall -u update -S interact
```

*(On the first node, upgrade data services from the Sun Cluster 3.1 Agents CD-ROM)*

```
phys-schost-1# ./scinstall -u update -s all -d /cdrom/cdrom0
```

*(On the second node, upgrade framework software from the Sun Cluster 3.1 CD-ROM)*

```
phys-schost-2# cd /cdrom/suncluster_3_1/SunCluster_3.1/Sol_8/Tools
phys-schost-2# ./scinstall -u update -S interact
```

*(On the second node, upgrade data services from the Sun Cluster 3.1 Agents CD-ROM)*

```
phys-schost-2# ./scinstall -u update -s all -d /cdrom/cdrom0
```

*(Reboot each node into the cluster)*

```
phys-schost-1# reboot
phys-schost-2# reboot
```

*(Verify cluster membership)*

```
scstat
-- Cluster Nodes --
 Node name Status
 -
Cluster node: phys-schost-1 Online
Cluster node: phys-schost-2 Online
```

## ▼ How to Upgrade Sun Cluster-Module Software for Sun Management Center

Perform the following steps to upgrade to the Sun Cluster 3.1 module software packages for Sun Management Center on the Sun Management Center server machine and help-server machine.

1. **Ensure that all Sun Management Center core packages are installed on the appropriate machines, as described in your Sun Management Center installation documentation.**

This step includes installing Sun Management Center agent packages on each cluster node.

2. **Become superuser on the Sun Management Center server machine.**
3. **Insert the Sun Cluster 3.1 CD-ROM into the CD-ROM drive.**
4. **Change to the**  
**/cdrom/suncluster\_3\_1/SunCluster\_3.1/Sol\_*ver*/Packages directory,**  
**where *ver* is 8 (for Solaris 8) or 9 (for Solaris 9) .**

```
cd /cdrom/suncluster_3_1/SunCluster_3.1/Sol_ver/Packages
```

5. **Install the Sun Cluster-module server package SUNWscssv.**  

```
pkgadd -d . SUNWscssv
```
6. **Change to the CD-ROM root directory and eject the CD-ROM.**
7. **Become superuser on the Sun Management Center help-server machine.**
8. **Repeat Step 3 through Step 6 to install the Sun Cluster-module help-server package SUNWscshl.**
9. **Finish the upgrade.**



Go to “How to Finish Upgrading to Sun Cluster 3.1 Software” on page 161.

## ▼ How to Finish Upgrading to Sun Cluster 3.1 Software

Perform this procedure to reregister and reversion all resource types that received a new version from the upgrade, then to re-enable resources and bring resource groups back online.

---

**Note** – To upgrade future versions of resource types, see “Upgrading a Resource Type” in *Sun Cluster 3.1 Data Service 4/03 Planning and Administration Guide*.

---

1. **Ensure that all steps in “How to Upgrade to Sun Cluster 3.1 Software” on page 154 are completed.**

2. **From any node, start the `scsetup(1M)` utility.**

```
scsetup
```

3. **To work with resource groups, type 2 (Resource groups).**

4. **To register resource types, type 4 (Resource type registration).**

Type **yes** when prompted to continue.

5. **Type 1 (Register all resource types which are not yet registered).**

The `scsetup` utility displays all resource types that are not registered.

Type **yes** to continue to register these resource types.

6. **Type 8 (Change properties of a resource).**

Type **yes** to continue.

7. **Type 3 (Manage resource versioning).**

Type **yes** to continue.

8. **Type 1 (Show versioning status).**

The `scsetup` utility displays which resources you can upgrade to new versions of the same resource type. The utility also displays the state that the resource should be in before the upgrade can begin.

Type **yes** to continue.

9. **Type 4 (Re-version all eligible resources).**

Type **yes** to continue when prompted.

10. **Return to the Resource Group Menu.**
11. **Type 6 (Enable/Disable a resource).**  
Type **yes** to continue when prompted.
12. **Select a resource to enable and follow the prompts.**
13. **Repeat Step 12 for each disabled resource.**
14. **When all resources are re-enabled, type q to return to the Resource Group Menu.**
15. **Type 5 (Online/Offline or Switchover a resource group).**  
Type **yes** to continue when prompted.
16. **Follow the prompts to bring each resource group online.**
17. **Exit the `scsetup` utility.**  
Type **q** to back out of each submenu, or press Ctrl-C.

The cluster upgrade is complete. You can now return the cluster to production.

---

## Recovering From Storage Configuration Changes During Upgrade

This section provides the following repair procedures to follow if changes were inadvertently made to the storage configuration during upgrade:

- “How to Handle Storage Reconfiguration During an Upgrade” on page 162
- “How to Resolve Mistaken Storage Changes During an Upgrade” on page 163

### ▼ How to Handle Storage Reconfiguration During an Upgrade

Any changes to the storage topology, including running Sun Cluster commands, should be completed before you upgrade the cluster to Solaris 9 software. If, however, changes were made to the storage topology during the upgrade, perform the following procedure. This procedure ensures that the new storage configuration is correct and that existing storage that was not reconfigured is not mistakenly altered.

1. **Ensure that the storage topology is correct.**

Check whether the devices that were flagged as possibly being replaced map to devices that actually were replaced. If the devices were not replaced, check for and correct possible accidental configuration changes, such as incorrect cabling.

**2. Become superuser on a node that is attached to the unverified device.**

**3. Manually update the unverified device.**

```
sddidadm -R device

-R device Performs repair procedures on the specified device
See the sddidadm(1M) man page for more information.
```

**4. Update the DID driver.**

```
sddidadm -ui
sddidadm -r

-u Loads the device ID configuration table into the kernel
-i Initializes the DID driver
-r Reconfigures the database
```

**5. Repeat Step 2 through Step 4 on all other nodes that are attached to the unverified device.**

**6. Return to the remaining upgrade tasks.**

Go to Step 11 in “How to Upgrade to Sun Cluster 3.1 Software” on page 154.

## ▼ How to Resolve Mistaken Storage Changes During an Upgrade

If accidental changes are made to the storage cabling during the upgrade, perform the following procedure to change the storage configuration back to the correct state.

---

**Note** – This procedure assumes that no physical storage was actually changed. If physical or logical storage devices were changed or replaced, instead follow procedures in “How to Handle Storage Reconfiguration During an Upgrade” on page 162.

---

**1. Change the storage topology back to its original configuration.**

Check the configuration of the devices that were flagged as possibly being replaced, including the cabling.

**2. As superuser, update the DID driver on each node of the cluster.**

```
scdidadm -ui
scdidadm -r

-u Loads the device-ID configuration table into the kernel
-i Initializes the DID driver
-r Reconfigures the database
```

See the `scdidadm(1M)` man page for more information.

**3. Did the `scdidadm` command return any error messages in Step 2?**

- If no, proceed to Step 4.
- If yes, return to Step 1 to make further modifications to correct the storage configuration, then repeat Step 2.

**4. Return to the remaining upgrade tasks.**

Go to Step 11 in “How to Upgrade to Sun Cluster 3.1 Software” on page 154.