



Sun Cluster 3.0 Release Notes

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900
U.S.A. 650-960-1300

Part Number 806-1428
November 2000, Revision A

Copyright Copyright 2000 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, California 94303-4900 U.S.A. All rights reserved.

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

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

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, OpenBoot, Solaris, Solstice DiskSuite, Sun Cluster, Sun Quad FastEthernet, Sun Management Center, Sun StorEdge, SunVTS, and Ultra 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.

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 2000 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, CA 94303-4900 Etats-Unis. Tous droits réservés.

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

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

Sun, Sun Microsystems, le Sun logo, AnswerBook2, docs.sun.com, OpenBoot, Solaris, Solstice DiskSuite, Sun Cluster, Sun Quad FastEthernet, Sun Management Center, Sun StorEdge, SunVTS, et Ultra 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.

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.

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



Contents

1. **Sun Cluster 3.0 Release Notes** 7
 - New Features 7
 - Supported Products 8
 - Installing Sun Cluster AnswerBooks 9
 - Setting Up the AnswerBook2 Documentation Server 10
 - Viewing Sun Cluster AnswerBooks 10
 - ▼ How to Install the Sun Cluster AnswerBooks 10
 - Viewing PDF Files 12
 - Sun Cluster 3.0 Restrictions 13
 - Supported Solaris Release and Patch Information 15
 - System Administration and Procedure Updates 16
 - `syncdir` Option Changes 16
 - Private Hostnames 16
 - Known Problems 16
 - Bug ID 4314698 16
 - Bug ID 4346123 17
 - Bug ID 4358349 17
 - Bug ID 4358629 18
 - Bug ID 4359321 19

Bug ID 4362435	19
Bug ID 4362925	19
Bug ID 4365310	20
Bug ID 4365700	20
Bug ID 4365729	21
Bug ID 4366840	21
Bug ID 4366886	21
Bug ID 4368034	22
Bug ID 4369228	22
Bug ID 4369565	22
Bug ID 4369668	23
Bug ID 4370760	23
Bug ID 4371236	24
Bug ID 4372369	24
Bug ID 4373498	25
Bug ID 4373911	25
Bug ID 4374194	25
Bug ID 4374648	26
Bug ID 4376171	26
Bug ID 4377303	26
Bug ID 4378553	26
Scalable Services Sticky Load-Balancing Policy	27
Upgrading Sun Cluster HA for Oracle from Sun Cluster 2.2 to Sun Cluster 3.0 Software	27
Known Documentation Problems	32
Known Problems With the Sun Management Center GUI	40
A. Sun Cluster Installation Configuration Worksheets and Examples	43
Installation Configuration Worksheets	43

Cluster and Node Names Worksheet	45
Cluster Interconnect Worksheet	46
Public Networks Worksheet	47
Local Devices Worksheet	48
Local File System Layout Worksheet	49
Disk Device Group Configurations Worksheet	50
Volume Manager Configurations Worksheet	51
Metadevices Worksheet (Solstice DiskSuite)	52
Installation Configuration Worksheet Examples	53
Example: Cluster and Node Names	54
Example: Cluster Interconnect	55
Example: Public Networks	56
Example: Local Devices	57
Example: Local File System Layout—With Mirrored Root	58
Example: Local File System Layout—Without Mirrored Root	59
Example: Disk Device Group Configurations	60
Example: Volume Manager Configurations	61
Example: Metadevices (Solstice DiskSuite)	62
B. Data Service Configuration Worksheets and Examples	63
Configuration Worksheets	63
Resource Types	64
Resource Groups—Failover	65
Resource Groups—Scalable	66
Network Resources	67
Application Resources	68
Configuration Worksheet Examples	69
Example: Resource Types	70
Example: Resource Groups—Scalable	71

Example: Resource Groups—Failover 72

Example: Network Resources—LogicalHostname 73

Example: Resources—SharedAddress 74

Example: Resources—Application 75

Sun Cluster 3.0 Release Notes

This document provides the following information for the Sun™ Cluster 3.0 software release:

- “New Features” on page 7
- “Supported Products” on page 8
- “Installing Sun Cluster AnswerBooks” on page 9
- “Viewing PDF Files” on page 12
- “Sun Cluster 3.0 Restrictions” on page 13
- “Supported Solaris Release and Patch Information” on page 15
- “System Administration and Procedure Updates” on page 16
- “Known Problems” on page 16

The appendices to this document include installation planning worksheets and examples to use when planning the Sun Cluster 3.0 software and data services installation. The Sun Cluster 3.0 AnswerBooks™ also include these worksheets.

New Features

This release includes the following new features:

- Integration with the Solaris™ operating environment kernel
- Cluster file system
- Global cluster-wide networking and device access
- Data services for scalable, multi-node HTTP with load balancing
- New, expanded API for developing highly available data services

- Ability to restart applications in parallel (Resource Group Manager)
- Failover support for applications that do not use shared disks
- Solaris 8 operating environment support

Supported Products

This section describes the supported software and memory requirements for Sun Cluster 3.0.

- **Operating environment and patches** – Supported Solaris versions and patch information are available through the following web site: <http://sunsolve.ebay.sun.com>. See the section “Supported Solaris Release and Patch Information” on page 15 for more details.
- **Volume managers** – Solstice DiskSuite™ 4.2.1 and VERITAS Volume Manager 3.0.4.
- **Data services (agents)** – Table 1–1, “Supported Data Services for Sun Cluster 3.0,” lists the supported data service versions and the Sun Cluster resource type names. Specify the resource type names when installing the data services using the `scinstall(1M)` utility and when registering the resource types associated with the data service using the `scradm(1M)` utility.

TABLE 1–1 Supported Data Services for Sun Cluster 3.0

Data Service	Sun Cluster Resource Type Name	Application Version
Sun Cluster 3.0 HA for DNS	<code>dns</code>	DNS on Solaris 8
Sun Cluster HA for NFS	<code>nfs</code>	NFS v2,3 on Solaris 8
Sun Cluster HA for iPlanet Web Server	<code>iws</code>	iPlanet Web Server 4.1
Sun Cluster 3.0 HA for Netscape Directory Server	<code>nsldap</code>	Netscape Directory Server 4.11
Sun Cluster 3.0 HA for Oracle	<code>oracle</code>	ORACLE 8.1.6 (8i)

TABLE 1-1 Supported Data Services for Sun Cluster 3.0 (continued)

Data Service	Sun Cluster Resource Type Name	Application Version
Sun Cluster HA for Oracle Parallel Server	N/A	ORACLE 8.1.6 (8i)
Sun Cluster 3.0 HA for Apache	apache	Apache HTTP 1.39

Note - Chapter 5, “Installing and Configuring Sun Cluster HA for Apache,” in the *Sun Cluster 3.0 Data Services Installation and Configuration Guide* describes the procedure for installing the Apache Web Server from the Apache web site. However, you can also install the Apache Web Server from the Solaris 8 operating environment CD-ROM. See the section “Data Services Installation and Configuration Guide” on page 38 of this document for more information.

- **Memory Requirements** – Sun Cluster 3.0 requires extra memory above what is configured for a node under a normal workload. The extra memory equals 128 Mbytes + 10%. For example, if a standalone node normally requires 1 Gbyte of memory, you need an extra 256 Mbytes to meet memory requirements.

Installing Sun Cluster AnswerBooks

The Sun Cluster 3.0 user documentation is available in AnswerBook2 format for use with AnswerBook2 documentation servers. The Sun Cluster 3.0 AnswerBook2 documentation set consists of:

- The Sun Cluster 3.0 Collection, which includes the following books:
 - Sun Cluster 3.0 Installation Guide*
 - Sun Cluster 3.0 System Administration Guide*
 - Sun Cluster 3.0 Hardware Guide*
 - Sun Cluster 3.0 Data Services Developers’ Guide*
 - Sun Cluster 3.0 Concepts*
 - Sun Cluster 3.0 Error Messages Manual*
- The Sun Cluster 3.0 Data Services Collection, which includes the following book:
 - Sun Cluster 3.0 Data Services Installation and Configuration Guide*

Setting Up the AnswerBook2 Documentation Server

The Solaris operating environment release includes AnswerBook2 documentation server software. The Solaris documentation CD-ROM, which is separate from the Solaris operating environment CD-ROM, includes the documentation server software. You need the Solaris documentation CD-ROM to install an AnswerBook2 documentation server.

If you have an AnswerBook2 documentation server installed at your site, you can use the same server for the Sun Cluster 3.0 AnswerBooks. If you do not have an AnswerBook2 documentation server installed, install a documentation server on a machine at your site. The administrative console that you use as the administrative interface to your cluster is a good choice for the documentation server. Do not use a cluster node as your AnswerBook2 documentation server.

For complete information on installing an AnswerBook2 documentation server, load the Solaris documentation CD-ROM on a server, and view the README files.

Viewing Sun Cluster AnswerBooks

Use the following procedure to view Sun Cluster 3.0 AnswerBooks from your AnswerBook2 documentation server. Install the Sun Cluster AnswerBook2 documents on a file system on the same server on which you install the documentation server. The Sun Cluster 3.0 AnswerBooks include a post-install script that automatically adds the documents to your existing AnswerBook library.

To use this procedure, you need the following:

- Root access to the server that has the AnswerBook2 documentation server installed. This server must have a CD-ROM drive. If you do not have an AnswerBook2 documentation server installed, you need the Solaris operating environment documentation CD-ROM. The Solaris operating environment CD-ROM contains the software and instructions you need to install an AnswerBook2 documentation server.
- The Sun Cluster 3.0 CD-ROMs that include the AnswerBook collections you are installing. The Sun Cluster 3.0 CD-ROM includes the Sun Cluster 3.0 Collection, and the Sun Cluster 3.0 Data Services CD-ROM contains the Sun Cluster 3.0 Data Services Collection.

▼ How to Install the Sun Cluster AnswerBooks

Use this procedure to install the Sun Cluster AnswerBook packages for the Sun Cluster 3.0 Collection and Sun Cluster 3.0 Data Services Collection.

- 1. Become superuser on the server that has the AnswerBook2 documentation server installed.**

2. **If you have previously installed the Sun Cluster AnswerBooks, remove the old packages.**

If you have never installed Sun Cluster AnswerBooks, ignore this step.

```
# pkgrm SUNWscfab SUNWscdab
```

3. **Insert the Sun Cluster CD-ROM or Sun Cluster Data Services CD-ROM into a CD-ROM drive attached to your documentation server.**

The Volume Management daemon, `vold(1M)`, should mount the CD-ROM automatically.

4. **Change directory to the location on the CD-ROM that contains the Sun Cluster AnswerBook package to install.**

The following directory contains the package for the Sun Cluster CD-ROM:
`suncluster_3_0/SunCluster_3.0/Packages.`

The following directory contains the package for the Sun Cluster Data Services CD-ROM: `sclatasesvices_3_0/components/SunCluster_Data_Service_Answer_Book_3.0/Packages.`

5. **Use the `pkgadd(1)` command to install the package.**

```
# pkgadd -d .
```

6. **Select the packages to install.**

Select the Sun Cluster 3.0 Collection (`SUNWscfab`) and the Sun Cluster 3.0 Data Services Collection (`SUNWscdab`).

7. **From the `pkgadd` installation options menu, choose `heavy` to add the complete package to the system and to update the AnswerBook2 catalog.**

Select either the Sun Cluster 3.0 Collection (`SUNWscfab`) or the Sun Cluster 3.0 Data Services Collection (`SUNWscdab`).

The document collection package included on each CD-ROM includes a post-install script that adds the collection to the documentation server's database and restarts the server. You should now be able to view the Sun Cluster AnswerBooks using your documentation server.

Viewing PDF Files

The Sun Cluster CD-ROMs now include a PDF file for each book in the Sun Cluster documentation set.

The following directory on the Sun Cluster CD-ROM contains the PDF files: `./suncluster_3_0/SunCluster_3.0/Docs/locale/C/PDF`.

The following directory on the Data Services CD-ROM contains the PDF file: `./scdataservices_3_0/components/SunCluster_Data_Service_Answer_Book_3.0/Docs/locale/C/PDF`.

Similar to the way the Sun Cluster AnswerBooks are delivered, six PDF files are being delivered on the Sun Cluster CD-ROM, and one PDF file is being delivered on the Data Services CD-ROM. Each PDF file is named with an abbreviation for the book the file contains.

Table 1-2, “Mapping of PDF Abbreviations to Book Titles,” shows the mapping of PDF file name abbreviations to the book titles.

TABLE 1-2 Mapping of PDF Abbreviations to Book Titles

CD-ROM	PDF Abbreviation	Book Title
Sun Cluster	CLUSTINSTALL	<i>Sun Cluster 3.0 Installation Guide</i>
	CLUSTNETHW	<i>Sun Cluster 3.0 Hardware Guide</i>
	CLUSTAPIPG	<i>Sun Cluster 3.0 Data Services Developers' Guide</i>
	CLUSTSYSADMIN	<i>Sun Cluster 3.0 System Administration Guide</i>
	CLUSTCONCEPTS	<i>Sun Cluster 3.0 Concepts</i>
	CLUSTERRMSG	<i>Sun Cluster 3.0 Error Messages Manual</i>
Data Services	CLUSTDATASVC	<i>Sun Cluster 3.0 Data Services Installation and Configuration Guide</i>

Sun Cluster 3.0 Restrictions

The following restrictions apply to the Sun Cluster 3.0 release:

- **Remote Shared Memory (RSM) transport types** – These transport types currently are not supported even though the documentation contains references to them.
- **Scalable Coherent Interface (SCI) is not supported as a cluster interconnect.**
- **The combination of VERITAS Volume Manager (VxVM) and Oracle Parallel Server (OPS) currently is not supported** – Use the hardware RAID volume manager with the Sun StorEdge™ A3500 to support OPS.
- **Automatic disk path monitoring is not supported** – You must manually monitor disk paths to make sure that a double failure or loss of path to a quorum device does not occur. The software detects active disk path failures, but no monitors detect inactive disk paths.
- **Storage devices with more than two physical paths to the enclosure are not supported** – The Sun StorEdge A3500, for which two paths are supported to each of two nodes, is an exception.
- **SunVTS™ is not supported.**
- **Only the Sun Cluster HA for NFS data service is supported for upgrade when using the `scinstall -u` command** – Upgrades of the framework from Sun Cluster 2.2 to Sun Cluster 3.0 software are fully supported.
- **Upgrade from Sun Cluster 2.2 to Sun Cluster 3.0 software is only supported for two-node clusters.**
- **Multihost tape and CD-ROM are not supported.**
- **Sun Cluster 3.0 software does not support the use of the loopback file system (LOFS) on cluster nodes.**
- **Do not run client applications on the cluster nodes** – Switchover or failover of a resource group might cause a TCP (`telnet/rlogin`) connection to be broken. This switchover or failover includes connections that the cluster nodes initiated and connections that client hosts outside the cluster initiated.
- **Do not run, on any cluster node, any processes that run in the time-sharing scheduling class with a higher-than-normal priority or any processes that run in the real-time scheduling class** – Sun Cluster 3.0 software relies on kernel threads that do not run in the real-time scheduling class. Other time-sharing processes that run at higher-than-normal priority or real-time processes could prevent the Sun Cluster kernel threads from getting needed CPU cycles and could cause various problems.
- **File system quotas are not supported in Sun Cluster 3.0 software.**
- **Logical network interfaces are reserved for use by Sun Cluster 3.0 software.**

- **The Sun Cluster 3.0 cluster file system does not support forced unmounts** – The command `umount -f` behaves the same as the `umount` command without the `-f` option.
- **Network Adapter Failover (NAFO) Restrictions**
 - All public networking adapters must be in NAFO groups. Sun Cluster 3.0 does not support public network adapters that are not in NAFO groups.
 - Only one NAFO group exists per IP subnet for each node. Sun Cluster 3.0 does not support even the weak form of IP striping, in which multiple IP addresses exist on the same subnet.
 - Only one adapter in a NAFO group can be active at any time.
 - Sun Cluster 3.0 does not support setting `local-mac-address?=true` in the OpenBoot™ PROM.
- **Service and Application Restrictions**
 - The Sun Cluster 3.0 product can be used to provide service for only those data services that are either supplied with the Sun Cluster product or set up using the Sun Cluster data services API.
 - Do not use cluster nodes as mail servers because the Sun Cluster environment does not support the `sendmail(1M)` subsystem. No mail directories should reside on Sun Cluster nodes.
 - Do not configure cluster nodes as routers (gateways). If the system goes down, the clients cannot find an alternate router and cannot recover.
 - Do not configure cluster nodes as NIS or NIS+ servers. Cluster nodes can, however, be NIS or NIS+ clients.
 - Do not use a Sun Cluster configuration to provide a highly available boot or install service on client systems.
 - Do not use a Sun Cluster 3.0 configuration to provide an `rarpd` service.
 - The Sun Cluster 3.0 data services API supports only 32-bit data services. The application on which the Sun Cluster data service depends can be a 64-bit application, but the data services' methods and monitors to support the application in a cluster must be 32-bit programs.
- **Sun Cluster 3.0 HA for NFS Restrictions**
 - Do not run, on any cluster node, any application that accesses a Sun Cluster HA for NFS file system on any other node. This intra-cluster use of Sun Cluster HA for NFS is not supported. Access such file systems only through the cluster file system. Using an NFS exported file system from a cluster node might lead to unpredictable locking behavior.
 - Sun Cluster HA for NFS requires that all NFS client mounts be “hard” mounts.
 - For Sun Cluster HA for NFS, do not use hostname aliases for network resources. NFS clients mounting cluster file systems using hostname aliases might experience `statd` lock recovery problems.

- Sun Cluster 3.0 does not support Secure NFS or the use of Kerberos with NFS. In particular, the `secure` and `kerberos` options to the `share_nfs(1M)` subsystem are not supported.
- **Volume Manager Restrictions**
 - In Solstice DiskSuite configurations using mediators, the number of mediator hosts configured for a diskset must be exactly two.
 - VxVM Dynamic Multipathing (DMP) with Sun Cluster 3.0 software is not supported.
 - Software RAID 5 is not supported.
- **Hardware Restrictions**
 - With the exception of clusters using Sun StorEdge A3x00, a pair of cluster nodes must have at least two multihost disk enclosures.
 - RAID 5 is only supported in hardware with the Sun StorEdge A3x00.
 - Alternate Pathing (AP) is not supported in Sun Cluster 3.0 configurations.
 - Gigabit Ethernet is supported for the cluster interconnect but not as a public network interface.

Supported Solaris Release and Patch Information

Access the SunSolve web pages at <http://sunsolve.ebay.sun.com> for the list of supported versions of the Solaris operating environment and required patches for Sun Cluster 3.0. Locate the Sun Cluster pages by doing a simple search specifying the `EarlyNotifier` collection and the search criteria “Sun Cluster 3.0.”

Review the `EarlyNotifier` information before installing Sun Cluster 3.0 and before you apply any patch to a cluster component (Solaris operating environment, Sun Cluster, volume manager, or disk firmware). All cluster member nodes must be at the same patch level for proper cluster operation.

Refer to the *Sun Cluster 3.0 System Administration Guide* for specific patch procedures and tips on administering patches.

System Administration and Procedure Updates

This section describes changes and updates to procedures used to administer a cluster.

syncdir Option Changes

In the Beta releases, you were required to specify the `syncdir` option when adding a cluster file system in `/etc/vfstab`. The GA release does not require this specification. Refer to the *Sun Cluster 3.0 Installation Guide* or the *Sun Cluster 3.0 Concepts* document for more information about this change.

Private Hostnames

Do not use the `scsetup` utility to change private hostnames after you have configured and started data services. Even though the `scsetup` utility permits you to change private hostnames, do not attempt to do so without contacting your Sun service representative.

Known Problems

The following known problems affect the operation of the Sun Cluster 3.0 GA release. The most current information on known problems can be accessed through the online *Release Notes* at <http://suncluster.eng.sun.com>.

Bug ID 4314698

Problem Summary: After installing Solstice DiskSuite software, the `scgdevs(1M)` command must be run for the Solstice DiskSuite device links to appear in the global namespace.

Workaround: Run the `scgdevs` command manually to make sure that Solstice DiskSuite device nodes are created.

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.

```
# fsck -y /dev/global/rdisk/d1s7
** /dev/global/rdisk/d1s7
Can't roll the log for /dev/global/rdisk/d1s7
```

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. Note that if any users of the file system exist on any of the cluster 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
```

Also, run the `share(1M)` command to confirm that the file system is not NFS shared on any of the cluster nodes.

Bug ID 4358349

Problem Summary: Do not create Sun Cluster HA for NFS resources in a resource group that contains a `SharedAddress` resource. Sun Cluster software does not support the use of `SharedAddress` resources with that data service.

Workaround: Add the desired logical hostname resources into the failover resource group.

You must set up a `LogicalHostname` resource with this step. The hostname used with Sun Cluster HA for NFS cannot be a `SharedAddress` resource.

```
# scrgadm -a -L -g resource-group-name -l hostname,...
```

-a -L -g resource-group-name Specifies the failover resource group into which to place the logical hostname resources.

-l hostname,... Specifies the network resources (logical hostnames) to be added.

Bug ID 4358629

Problem Summary: Upgrades from Sun Cluster 2.2 to Sun Cluster 3.0 software might fail if the logical hosts created for the Sun Cluster 2.2 software use a number for the IP address rather than a hostname.

Workaround: The two ways to solve this problem are:

- If the cluster is running Sun Cluster 2.2 software, rerun the `scconf(1M)` command to reset the LOGIP field to an actual name instead of an IP address.

```
# scconf clustname -L logical_host -n host_list -g diskgroup_list -i adapter_list,logical_hostname
```

The following example shows how to reset the LOGIP field to an actual name instead of an IP address.

```
# scconf schost -L schost-1 -n phys-schost-1,phys-schost-2 -g schost-1 -i hme0,hme0,schost-1
```

- If the cluster has already been upgraded to Sun Cluster 3.0 software, edit the saved copy of the `ccd` file on the upgraded nodes. Edit the `ccd` file in the `/var/cluster/logs/install/preserve/2.2/SUNWcluster/conf` directory, replacing the LOGIP field with an actual name.

For example, the IP address specified in the following line should be changed to the IP address' corresponding logical hostname.

```
LOGIP:phys-schost-1,phys-schost-2:hme0,hme0:192.29.72.238:1
```

The corresponding logical hostname replaces the IP address in the `ccd` file.

Bug ID 4359321

Problem Summary: The `scinstall` utility enables you to specify the `/global` directory for the global devices file system. However, because the mount point for the global devices file system is `/global/.devices/node@nodeid`, this specification should not be enabled.

Workaround: Re-install the node using the correct name for the global devices file system.

Although not preferred, fixing the entries in the `/etc/vfstab` files, rebooting the cluster, and then running the `scgdevs` command is a possible workaround. Check that each `/global/.devices/node@nodeid` entry in each `/etc/vfstab` file has the global mount option set.

Bug ID 4362435

Problem Summary: When the Sun Cluster 3.0 module is loaded into the Sun Management Center 2.1 console, and you try to access Resource Type Definition->Properties Table, if the table is more than one page long, it never loads.

Workaround: Run the `scrgadm -pvv` command to see all resource type properties.

Bug ID 4362925

Problem Summary:

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

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

```
nodeA# umount /global/oracle
umount: global/oracle busy
nodeA# fuser -c /global/oracle
/global/oracle: nodeA# umount /global/oracle
umount: global/oracle busy
nodeA# fuser -c /global/oracle
/global/oracle: 335co 317co 302co 273co 272co
nodeA# ps -ef|grep 335
```

(continued)

```

root 335 273 0 17:17:41 ?          0:00 /usr/sbin/nsrmmmd -n 1
root 448 397 0 17:19:37 console 0:00 grep 335

```

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

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

Bug ID 4365310

Problem Summary: If a resource state becomes `STOP_FAILED`, you must manually clear the `STOP_FAILED` flag for the resource. If you specify multiple resources to have their flags cleared, and one of the resources is not in the `STOP_FAILED` state, the function returns early without clearing the `STOP_FAILED` flags of the other resources listed.

No error message is displayed in this case, but the flags for the other resources are not cleared. The lack of an error message is misleading, giving no indication that a failure occurred, while the `STOP_FAILED` state is not cleared for all of the resources listed in the command.

Workaround: To avoid this problem, clear the `STOP_FAILED` flags individually for each resource that are in `STOP_FAILED` state.

```
# scswitch -c -f STOP_FAILED -j stopfailres -h phys-schost-1
```

Bug ID 4365700

Problem Summary: In the following example, multiple resources are disabled from the same resource group with a single command.

```
# scswitch -n -j r1,r2,r3
```

If the first resource moves into the `STOP_FAILED` state, the remaining resources might end up disabled but still online. This online state represents an invalid internal state of the Resource Group Manager daemon and can cause the Resource Group Manager daemon to panic.

Workaround: When disabling resources, always disable just one resource per `scswitch(1M)` command.

Bug ID 4365729

Problem Summary: Attempts to put a device group into maintenance mode using the following command fail if file systems are mounted on the specified device group.

```
# scswitch -m -D device-group
```

Workaround: Unmount all file systems on the device group to be put into maintenance. A device group can be placed in maintenance mode only if the devices in that device group are unused, meaning no active users of devices exist in that device group, and all dependent filesystems are unmounted.

Bug ID 4366840

Problem Summary: If any cables and associated adapters or junctions are removed from a cluster while one of its nodes is down, that node will panic when it is rebooted and attempts to rejoin the cluster.

Workaround: Until this bug is fixed, do not remove cables, adapters, or junctions from a cluster while a node is down. If you do experience this panic, reboot the node a second time. The node can then join the cluster without panicking.

Bug ID 4366886

Problem Summary: Heavy system load might interfere with bringing device groups online. This problem occurs because VERITAS Volume Manager (VxVM) needs to perform several tasks, such as syncing mirrors, to import a disk group. Under heavy load, these tasks can be prevented from completing in a timely manner because other system tasks are utilizing important system resources. Because device groups are commonly brought online automatically when a node boots (if a file system is set to automatically mount, for example), such an online hang might manifest itself as a hang during boot.

Workaround: Decrease system load or increase the priority of the vxconfigd daemon.

Bug ID 4368034

Problem Summary: If the Resource Group Manager daemon dies, or a node dies while a remote procedure call is in progress, error messages such as one of the following might be printed on the system console.

```
COMM_FAILURE SystemException: COMM_FAILURE major 3 minor 0 Error 0 completed NO
```

```
INV_OBJREF SystemException: INV_OBJREF major 4 minor 9 Bad file number completed NO
```

These messages are intended for debugging use rather than customer consumption. The Resource Group Manager daemon already writes clearer `syslog` messages for these exceptions, so the debugging `printf`'s are unnecessary.

Workaround: Ignore these console messages. Look for `syslog` messages regarding a node death. Normally, the Resource Group Manager daemon recovers automatically from such an event.

Bug ID 4369228

Problem Summary: The `dbassist` utility provided by Oracle does not enable creation of an Oracle Parallel Server database directly on a hardware RAID device.

Workaround: Use the Oracle Server Manager line mode, `svrgmr1`, to create Oracle Parallel Server databases on Sun Cluster 3.0 software.

Bug ID 4369565

Problem Summary: The `nfs_upgrade` script is not idempotent. You cannot run the script twice.

Workaround: If you need to run the script twice, remove the NFS resource and NFS resource type that were created in the first attempt before running the script a second time.

Bug ID 4369668

Problem Summary: When the system administrator edits the `Nodelist` property of a managed resource group, the Resource Group Manager should run the `INIT` method on all resources in the resource group that have the property `Init_nodes=RG_PRIMARYES`, on all nodes that have been added to the node list. The Resource Group Manager should run the `FINI` method on such resources, on nodes that were deleted from the node list. Similarly, if the `Installed_nodes` property of a resource type is edited, the Resource Group Manager should run the `INIT` or `FINI` method on all resources of that type that reside in managed resource groups and have the property `Init_nodes=RT_installed_nodes`.

Currently, the Resource Group Manager does not run `INIT` or `FINI` methods when these updates are performed. As a result, the resources might not be properly initialized or cleaned up on these nodes.

Workaround: Using the `scswitch` command, unmanage and then re-manage the affected resource groups. Unfortunately, this process requires that the administrator take the resource group offline. Alternatively, the administrator can run the equivalent `INIT` or `FINI` actions manually (without unmanaging the resource group), if such procedures are documented for the resource types that occur within the group.

This workaround is unnecessary if none of the resources in the group have `INIT` or `FINI` methods. The only Sun-supplied resource types that use `INIT` and `FINI` methods are:

- `SUNW.SharedAddress` (shared address resource type)
- `SUNW.apache` (Apache Web Server) configured in scalable mode
- `SUNW.iws` (iPlanet/Netscape web server) configured in scalable mode

Resource types that customers or third parties implement might also use `INIT` or `FINI` methods. If so, this workaround is necessary for resource groups that contain such resource types.

Note - All scalable services implicitly use `INIT` and `FINI` methods, even if such methods are not explicitly declared for the resource type.

Bug ID 4370760

Problem Summary: You cannot remove the last host from a metaset unless you first take the device group offline.

Workaround: To remove the last host from a metaset, take the device group offline first. To remove the last host, run the following two commands as superuser from the host to be removed.

```
# /usr/cluster/bin/scswitch -m -D disksetname
# metaset -s disksetname -d -h hostname
```

Bug ID 4371236

Problem Summary: Some `ge` switches require some of the `ge` device parameters to be set to values other than the default values. Chapter 3 of the *Sun GigabitEthernet/P 2.0 Adapter Installation and User's Guide* describes the procedure to change `ge` device parameters. The procedure to be used on nodes running Sun Cluster 3.0 software varies slightly from that described in the guide. In particular, the difference is in how the device path names in the `/etc/path_to_inst` file are used to derive parent names for use in the `ge.conf` file.

Workaround: Chapter 3 of the *Sun GigabitEthernet/P 2.0 Adapter Installation and User's Guide* describes the procedure to change `ge` device parameter values through entries in the `/kernel/drv/ge.conf` file. The procedure to determine the parent name from the `/etc/path_to_inst` listing (to be used in `ge.conf` entries) appears on page 24, "Setting Driver Parameters Using a `ge.conf` File." For example, from the following `/etc/path_to_inst` line, you can determine the parent name for `ge2` to be `/pci@4,4000`.

```
``/pci@4,4000/network@4'' 2 ``ge''
```

On cluster nodes, you must delete the `/node@nodeid` prefix in the device paths in `/etc/path_to_inst` prior to using that prefix as a parent name. For example, on a cluster node, an equivalent `/etc/path_to_inst` entry could have been the following entry.

```
``/node@1/pci@4,4000/network@4'' 2 ``ge''
```

The parent name for `ge2` to use in `ge.conf` is still `/pci@4,4000`.

Bug ID 4372369

Problem Summary: The `nfs_upgrade` script cannot work if more than one logical host is configured in Sun Cluster 2.2 software.

Workaround: No current workaround exists. If you encounter this problem, contact your Sun Service provider about acquiring a patch.

Bug ID 4373498

Problem Summary: The LDAP administrative server treats hostnames as case sensitive. While working with the LDAP administrative server, therefore, all hostnames specified in the LDAP configuration should match case with the LDAP specification in the name service in use on the cluster node. This case matching is particularly important if DNS is the name service in use because the DNS domain name must also match exactly with the hostname specification in the LDAP configuration.

Workaround: Make sure the case of the fully qualified domain name of the machine given to LDAP matches the case of the domain name returned by the resolver.

Bug ID 4373911

Problem Summary: If you do the following:

- Configure multiple Sun Cluster HA for NFS resources on the cluster.
 - Switch an HA-NFS resource on or off and leave another online on the same node.
- the HA-NFS fault monitor might display the following warning message.

```
clnt_tp_create_timed of program statd failed:RPC:Program not registered
```

Workaround: No workaround is necessary. The warning message can safely be ignored.

Bug ID 4374194

Problem Summary: The Sun Management Center agent might unexpectedly exit on Ultra™ 2 workstations with Sun StorEdge A5000. The problem occurs when the Sun Management Center agent is set up with Config Reader, and the Config-Reader4udt module is added to the `/var/opt/SUNWsymon/cfg/base-modules-d.dat` file. The Sun Management Center agent reads this file on startup and then tries to load all listed modules. The agent might segmentation fault while trying to load the Config-Reader4udt module.

Workaround: To avoid this problem, do one of the following:

- Delete Config-Reader4udt from the `base-modules-d.dat` file.
- Replace Config-Reader4udt with Config-Reader4u, which should provide the same information for Ultra 2 machines.
- At Sun Management Center agent setup time, do not set up the Config Reader module at all. This prevents the Config-Reader4udt module from being added to the `base-modules-d.dat` file.

Bug ID 4374648

Problem Summary: The `scinstall` man page currently has an example that uses `-s oracle` to automatically upgrade a Sun Cluster HA for Oracle data service from Sun Cluster 2.2 to Sun Cluster 3.0 software. This option is currently unsupported.

Workaround: Do not use the `-s oracle` option to attempt to upgrade from Sun Cluster 2.2 to Sun Cluster 3.0 software for an Oracle data service. Instead, use the manual upgrade procedure, “Upgrading Sun Cluster HA for Oracle from Sun Cluster 2.2 to Sun Cluster 3.0 Software” on page 27.

Bug ID 4376171

Problem Summary: Placing a FC-AL SBus Card (FC100/S) and a Sun Quad FastEthernet™ 2.0 (SQFE/S) on the same SBus might cause unexpected resets on the QFE card.

Workaround: Avoid configuring cluster nodes with a FC-AL SBus Card (FC100/S) and a Sun Quad FastEthernet 2.0 (SQFE/S) on the same SBus.

Bug ID 4377303

Problem Summary: Newly created Sun StorEdge A3500 LUNs might not always appear in format on all nodes.

Workaround: Run the `/etc/raid/bin/hot_add` command on nodes that do not see the new LUNs.

Bug ID 4378553

Problem Summary: A resource group’s `Nodelist` property is an ordered list of nodes that can master the resource group, with the most-preferred node listed first. The Resource Group Manager should always host a resource group on the most-preferred node that is available. However, when an administrator reboots the cluster (when all nodes are rebooting at once), managed resource groups might end up being mastered on nodes other than the most-preferred node. This problem occurs only upon reboot of the entire cluster.

Workaround: After rebooting the cluster, use the `scswitch` command to switch resource groups onto the desired nodes. The `Nodelist` preference order will be enforced automatically from that point onward, as long as the cluster remains up.

Scalable Services Sticky Load-Balancing Policy

Currently, you might encounter a problem if you run a scalable data service that uses the sticky load-balancing policy. The problem can occur if the service runs with stickiness established relative to a particular node, and later you start another instance of the same service on a different node. Starting another instance of the same service might cause the first instance to lose its stickiness.

The result that the sticky algorithm returns when the second instance starts determines whether the first instance loses its stickiness. The algorithm should not change the sticky affinity in this case, but sometimes the algorithm does change the sticky affinity.

Refer to *Sun Cluster 3.0 Concepts* for more information on the sticky load-balancing policy.

Upgrading Sun Cluster HA for Oracle from Sun Cluster 2.2 to Sun Cluster 3.0 Software

Perform these procedures while upgrading the Sun Cluster framework using the `scinstall upgrade` procedure.

Conditions and Restrictions

The following conditions and restrictions apply when upgrading Sun Cluster HA for Oracle from Sun Cluster 2.2 to Sun Cluster 3.0 software.

- Before upgrading Sun Cluster HA for Oracle, you must check that you have met all conditions described in the *Sun Cluster 3.0 Installation Guide* relating to upgrading the Sun Cluster framework from Sun Cluster 2.2.
- Sun Cluster HA for Oracle must be functioning normally on Sun Cluster 2.2 before you begin the upgrade.

How to Save the Sun Cluster HA for Oracle Configuration Files

Use the following procedure to save the configuration files from your Sun Cluster 2.2 configuration.

1. **Follow the `scinstall framework-upgrade` procedure until you have completed the upgrade-begin steps (`scinstall -F begin`) on each node.**
2. **Run the following command on each node as superuser. This command will save a version of all files in the `/var/opt/oracle` directory.**

To ensure that this information does not get lost, back up the structure found in the `/var/opt/oracle` directory to an external device.

```
# cp -r /var/opt/oracle /var/cluster/logs/install/preserve/2.2/SUNWscor
```

3. Complete the finish portion of the framework upgrade (`scinstall -u finish`).

Note - Do not use the `-s oracle` option with the `scinstall -u finish` command. This option attempts an automated upgrade for Sun Cluster HA for Oracle, and the automated upgrade will fail. The only automated upgrade supported is for NFS.

After completing the framework upgrade, set up the Sun Cluster 3.0 environment. The following section, “Setting Up the Sun Cluster 3.0 Environment” on page 28, describes this procedure.

Setting Up the Sun Cluster 3.0 Environment

Perform the following steps to set up your Sun Cluster 3.0 environment.

1. On one node, run the following command to verify that:

- The framework upgrade has correctly set up a Sun Cluster 3.0 resource group that corresponds to each Sun Cluster 2.2 logical host.
- The hostname network resource is in the resource group and is online.

```
# scstat -g
```

2. On one node, run the following command to verify that the VERITAS disk group or Solstice DiskSuite diskset that held the Oracle database (and possibly the Oracle binaries) in Sun Cluster 2.2 is correctly mapped into a Sun Cluster 3.0 disk device group.

```
# scstat -D
```

3. On each node, run the following command to verify that the required file systems for each Oracle instance are mounted.

```
# mount
```

4. On each node, run the following commands to restore the saved version of the Oracle configuration files under the `/var/opt` directory.

If you saved the files in the `/var/opt/oracle` directory earlier in the procedure, and the files are unchanged, you can skip this step.

```
# cp -r /var/cluster/logs/install/preserve/2.2/SUNWscor/oracle /var/opt
# chown -R oracle:dba /var/opt/oracle
```

Configure Sun Cluster HA for Oracle Under Sun Cluster 3.0

Configure Sun Cluster 3.0 HA for Oracle using the following procedure.

Note - Perform Step 1 only once.

1. On one node, register the Oracle server and listener resource types using the following commands.

```
# scrgadm -a -t SUNW.oracle_server
# scrgadm -a -t SUNW.oracle_listener
```

Run Step 2 on page 29 through Step 5 on page 31 for each Sun Cluster 2.2 HA for Oracle instance listed in the `/var/opt/oracle/oratab` file.

2. Determine the value of the `ORACLE_HOME` variable from the `oratab` file.

For example, suppose the `oratab` file shows the following information.

ora32:/oracle/816_32:N

This information indicates that the `ORACLE_HOME` variable for the `ORACLE_SID` `ora32` instance is the value `/oracle/816_32`.

3. Retrieve the parameter values from the `ccd.database` file for each Oracle instance.

These parameters will map into Sun Cluster 3.0 parameters to `scrgadm`. You will use these parameters when configuring Sun Cluster HA for Oracle under Sun Cluster 3.0.

```
# grep ^HAORACLE: /var/cluster/logs/install/preserve/2.2/SUNWcluster/conf/ccd.database
```

Each Oracle instance in the `ccd.database` file uses the following format

```
HAORACLE:on:ora32:boots-1:60:10:120:300:scott/tiger:/oracle/816_32/dbs/initora32.ora:ORA_LIST
```

These parameters map into the following Sun Cluster 3.0 format.

```
HAORACLE:STATE:ORACLE_SID:LOGICAL_HOSTNAME_IP_Resource:THOROUGH_PROBE_INTERVAL:CONNECT_CYCLE:PROBE_TIMEOUT:RETRY_INTERVAL
```

The resource group name `RG_NAME` will be

`${LOGICAL_HOSTNAME_IP_Resource}-lh`. Note that the `-lh` will be automatically appended to the resource group name in Sun Cluster 3.0.

4. Locate the `background_dump_dest` value in the `$PARAMETER_FILE` variable, and set the `ALERT_LOG_FILE` variable to the following value.

```
$background_dump_dest/alert_${ORACLE_SID}.log
```

For example, for ORACLE_SID=ora32, suppose that in the \$PARAMETER_FILE file, background_dump_dest is the following value.

```
/oracle/816_32/admin/ora32/bdump
```

In this example, ALERT_LOG_FILE should be updated to the following value.

```
/oracle/816_32/admin/ora32/bdump/alert_ora32.log
```

5. On one node, run the following commands to create Oracle resources and bring them online.

```
# scrgadm -a -t SUNW.oracle_server -g $RG_NAME -j $ORACLE_SID-serv \  
-x Oracle_sid=$ORACLE_SID -x Oracle_home=$ORACLE_HOME \  
  
-y Thorough_probe_interval=$THOROUGH_PROBE_INTERVAL \  
  
-x Connect_cycle=$CONNECT_CYCLE -x Probe_timeout=$PROBE_TIMEOUT \  
  
-y Retry_interval=$RETRY_INTERVAL -x Connect_string=$CONNECT_STRING \  
  
-x Parameter_file=$PARAMETER_FILE -x Alert_log_file=$ALERT_LOG_FILE  
# scrgadm -a -j $ORACLE_SID-list -t SUNW.oracle_listener -g $RG_name \  
-x Oracle_home=$ORACLE_HOME -x Listener_name=$LISTENER_NAME  
# scswitch -e -j $ORACLE_SID-serv  
# scswitch -e -j $ORACLE_SID-list  
# scswitch -e -M -j $ORACLE_SID-serv  
# scswitch -e -M -j $ORACLE_SID-list
```

For example, using the Oracle instance described in Step 2 on page 29, Step 3 on page 30, and Step 4 on page 30, you would run the following commands.

```
# scrgadm -a -t SUNW.oracle_server -g boots-1-lh -j ora32-serv \  
-x Oracle_sid=ora32 -x Oracle_home=/oracle/816_32 \  
  
-y Thorough_probe_interval=60 \  
  
-x Connect_cycle=10 -x Probe_timeout=120 \  
  
-y Retry_interval=300 -x Connect_string=scott/tiger \  
  
-x Parameter_file=/oracle/816_32/dbs/initora32.ora \  
  
-x Alert_log_file=/oracle/816_32/admin/ora32/bdump/alert_ora32.log  
# scrgadm -a -j ora32-list -t SUNW.oracle_listener -g boots-1-lh \  
-x Oracle_home=/oracle/816_32 -x Listener_name=ORA_LIST  
# scswitch -e -j ora32-serv
```

(continued)

```
# scswitch -e -j ora32-list
# scswitch -e -M -j ora32-serv
# scswitch -e -M -j ora32-list
```

Verify the Upgrade

To verify that the upgrade has completed successfully, perform the following steps.

1. **Verify that the Oracle resources are online by using the following command**

```
# scstat -g
```

.

2. **Verify that you can switch over the resource group by using the following command.**

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

Known Documentation Problems

This section discusses documentation errors you might encounter and steps to correct these problems.

Installation Guide

The *Sun Cluster 3.0 Installation Guide* contains the following documentation errors:

- In Step 11a of the procedure “How to Use JumpStart to Install the Solaris Operating Environment and Establish New Cluster Nodes,” the following command is incorrectly formatted.


```
# mount | grep global | egrep -v node@ | awk '{print $1}'
```

The command should instead read as follows.

```
# mount | grep global | egrep -v node@ | awk '{print $1}'
```

The two apostrophes (') in the command are the same character. They are not meant to represent open and closed single quotation marks.

- In the installation and upgrade procedures, the CD-ROM directory paths are incorrect. Whenever a procedure uses `/cdrom_image` as part of a CD-ROM directory path, replace that part of the path with `/cdrom`.

For example, in Step 3 of the procedure “How to Install Cluster Control Panel Software to the Administrative Console,” the following CD-ROM path is given.

```
# cd /cdrom_image/suncluster_3_0/SunCluster_3.0/Packages
```

Use the following directory path instead.

```
# cd /cdrom/suncluster_3_0/SunCluster_3.0/Packages
```

Hardware Guide

In the *Sun Cluster 3.0 Hardware Guide*, the following procedures are incorrect or do not exist:

- The task map “Configuring StorEdge A3500 Disk Drives” in Chapter 7 contains an error. The task “Increase the drive capacity of a LUN” is titled incorrectly. You cannot increase the drive capacity of a LUN. You can, however, increase the size of the drive group. Therefore, you do not need to remove the LUN from any diskset or disk group. The task should read: “Increase the size of the drive group. Follow the same procedure used in a non-cluster environment.”
- No procedures are documented in the Sun Cluster 3.0 GA AnswerBook for recabling disks without introducing redundant disk paths in the Cluster Configuration Repository (CCR).

When you recable devices within a cluster, the cluster must become aware of the new device configuration. To ensure that the cluster becomes aware of the new configuration and to guarantee device availability, use the following procedures:

- “How to Move a Disk Cable to a New Adapter” on page 34

- “How to Move a Disk Cable From One Node to Another” on page 34
- “How to Update Cluster Software to Reflect Proper Device Configuration” on page 36

How to Move a Disk Cable to a New Adapter

Use the following procedure to move a disk cable to a new adapter within a node.

1. **Quiesce all I/O to the affected disk(s).**
2. **Unplug the cable from the old adapter.**
3. **Run the `cfgadm(1M)` command on the local node to unconfigure all drives affected by the move.**

Or, reboot the node by using the following command.

```
# reboot -- -r
```

4. **Run the `devfsadm -C` command on the local node to clean up the Solaris device link.**
5. **Run the `scdidadm -C` command on the local node to clean up the DID device path.**
6. **Connect the cable to the new adapter.**
7. **Run the `cfgadm` command on the local node to configure the drives in the new location.**

Or, reboot the node by using the following command.

```
# reboot -- -r
```

8. **Run the `scgdevs` command to add the new DID device path.**

How to Move a Disk Cable From One Node to Another

Use the following procedure to move a disk cable from one node to another node.

1. Delete all references to the path you wish to remove from all volume manager and data service configurations.
2. Quiesce all I/O to the affected disk(s).
3. Unplug the cable from the old node.
4. Run the `cfgadm` command on the old node to unconfigure all drives affected by the move.
Or, reboot the node by using the following command.

```
# reboot -- -r
```

5. Run the `devfsadm -C` command on the old node to clean up the Solaris device link.
6. Run the `scdidadm -C` command on the old node to clean up the DID device path.
7. Connect the cable to the new node.
8. Run the `cfgadm` command on the new node to configure the drives in the new location.
Or, reboot the node by using the following command.

```
# reboot -- -r
```

9. Run the `devfsadm` command on the new node to create the new Solaris device links.
10. Run the `scgdevs` command on the new node to add the new DID device path.
11. Add the path on the new node to any required volume manager and data service configurations.
When configuring data services, check that your node failover preferences are set to reflect the new configuration.

How to Update Cluster Software to Reflect Proper Device Configuration

If the preceding procedures are not followed correctly, an error might be logged the next time you run the `scdidadm -r` command or the `scgdevs` command. To update the cluster software to reflect the proper device configuration, perform the following steps.

1. **Make sure cable configuration is as you want it to be. Make sure the cable is detached from the old node.**
2. **Make sure the old node is removed from any required volume manager or data service configurations.**
3. **Run the `cfgadm` command on the old node to unconfigure all drives affected by the move.**

Or, reboot the node by using the following command.

```
# reboot -- -r
```

4. **Run the `devfsadm -C` command on the node from where you removed the cable.**
5. **Run the `scdidadm -C` command on the node from where you removed the cable.**
6. **Run the `cfgadm` command on the new node to configure the drives in the new location.**

Or, reboot the node by using the following command.

```
# reboot -- -r
```

7. **Run the `scgdevs` command on the new node to add the new DID device path.**
8. **Run the `scdidadm -R device` command on the new node to make sure that SCSI reservations are in the correct state.**

Data Services Developers' Guide

The sample code in Appendix B of the *Sun Cluster 3.0 Data Services Developers' Guide* has two known problems:

- Many of the lines in the code listings wrap incorrectly, particularly in extended comments. To see the correct line wrapping, view the PDF version of the manual.
- A variable declaration is missing from most of the method script listings in this appendix. Each method should declare the following variable in `main()`:

```
SYSLOG_TAG=$RESOURCETYPE_NAME,$RESOURCEGROUP_NAME,$RESOURCE_NAME
```

This variable is used throughout the sample code in the `logger()` command.

Concepts Guide

The following points should be noted about *Sun Cluster 3.0 Concepts*:

- The diagrams throughout *Sun Cluster 3.0 Concepts* do not display correctly in AnswerBook. The text callouts did not get sized correctly when they were converted for use by AnswerBook. The PDF version of *Sun Cluster 3.0 Concepts*, which is included on the Sun Cluster CD-ROM, shows the callouts correctly.
- *Sun Cluster 3.0 Concepts* does not include the following section: “Using the Cluster Interconnect for Application Traffic” on page 37 of this document. This section explains how data service developers and system administrators can use the cluster interconnect for application traffic.

Using the Cluster Interconnect for Application Traffic

A cluster must have multiple network connections between nodes, forming the cluster interconnect. The clustering software uses multiple interconnects both for high availability and to improve performance. For internal traffic (for example, file system data or scalable services data), messages are striped across all available interconnects in a round-robin fashion.

The cluster interconnect is also available to applications, for highly available communication between nodes. For example, a distributed application might have components running on different nodes that need to communicate. By using the cluster interconnect rather than the public interconnect, these connections can withstand the failure of an individual link.

To use the cluster interconnect for communication between nodes, an application must use the private hostnames configured when the cluster was installed. For example, if the private hostname for node 1 is `clusternode1-priv`, use that name to communicate over the cluster interconnect to node 1. TCP sockets opened using

this name are routed over the cluster interconnect and can be transparently re-routed in the event of network failure.

Note that because the private hostnames can be configured during installation, the cluster interconnect can use any name chosen at that time. The actual name can be obtained from `scha_cluster_get(3HA)` with the `scha_privatelink_hostname_node` argument.

For application-level use of the cluster interconnect, a single interconnect is used between each pair of nodes, but separate interconnects are used for different node pairs if possible. For example, consider an application running on three nodes and communicating over the cluster interconnect. Communication between nodes 1 and 2 might take place on interface `hme0`, while communication between nodes 1 and 3 might take place on interface `qfe1`. That is, application communication between any two nodes is limited to a single interconnect, while internal clustering communication is striped over all interconnects.

Note that the application shares the interconnect with internal clustering traffic, so the bandwidth available to the application depends on the bandwidth used for other clustering traffic. In the event of a failure, internal traffic can round-robin over the remaining interconnects, while application connections on a failed interconnect can switch to a working interconnect.

Two types of addresses support the cluster interconnect, and `gethostbyname(3N)` on a private hostname normally returns two IP addresses. The first address is called the *logical pairwise address*, and the second address is called the *logical pernode address*.

A separate logical pairwise address is assigned to each pair of nodes. This small logical network supports failover of connections. Each node is also assigned a fixed pernode address. That is, the logical pairwise addresses for `clusternode1-priv` are different on each node, while the logical pernode address for `clusternode1-priv` is the same on each node. A node does not have a pairwise address to itself, however, so `gethostbyname(clusternode1-priv)` on node 1 returns only the logical pernode address.

Note that applications accepting connections over the cluster interconnect and then verifying the IP address for security reasons must check against all IP addresses returned from `gethostbyname`, not just the first IP address.

If you need consistent IP addresses in your application at all points, configure the application to bind to the pernode address on both the client and the server side so that all connections can appear to come and go from the pernode address.

Data Services Installation and Configuration Guide

Chapter 5, "Installing and Configuring Sun Cluster HA for Apache," in the *Sun Cluster 3.0 Data Services Installation and Configuration Guide* describes the procedure for installing the Apache Web Server from the Apache web site (<http://www.apache.org>). However, you can also install the Apache Web Server from the Solaris 8 operating environment CD-ROM.

The Apache binaries are included in three packages—SUNWapchr, SUNWapchu, and SUNWapchd—that form the SUNWCapache package metacluster. You must install SUNWapchr before SUNWapchu.

Place the Web server binaries on the local file system on each of your cluster nodes or on a cluster file system.

▼ Installing Apache from the Solaris 8 CD-ROM

This procedure documents the steps required to use the Sun Cluster HA for Apache data service with the version of the Apache Web Server that is on the Solaris 8 operating environment CD-ROM.

- 1. Install the Apache packages** SUNWapchr, SUNWapchu, **and** SUNWapchd **if they are not already installed.**

Use `pkginfo(1)` to determine if the packages are already installed.

```
# pkgadd -d Solaris 8 Product directory SUNWapchr SUNWapchu SUNWapchd
...
Installing Apache Web Server (root) as SUNWapchr
...
[ verifying class initd ]
/etc/rc0.d/K16apache linked pathname
/etc/rc1.d/K16apache linked pathname
/etc/rc2.d/K16apache linked pathname
/etc/rc3.d/S50apache linked pathname
/etc/rcS.d/K16apache linked pathname
...
```

- 2. Disable the start and stop run control scripts that were just installed as part of the SUNWapchr package.**

Disabling these scripts is necessary because the Sun Cluster HA for Apache data service will start and stop the Apache application after the data service has been configured. Perform the following steps:

- a. List the Apache run control scripts.**
- b. Rename the Apache run control scripts.**
- c. Verify that all the Apache-related scripts have been renamed.**

Note - The following example changes the first letter in the name of the run control script from upper case to lower case. You can rename the scripts, however, in a fashion consistent with your normal administration practices.

```
# ls -l /etc/rc?.d/*apache
/etc/rc0.d/K16apache
/etc/rc1.d/K16apache
/etc/rc2.d/K16apache
/etc/rc3.d/S50apache
/etc/rcS.d/K16apache

# mv /etc/rc0.d/K16apache /etc/rc0.d/k16apache# mv /etc/rc1.d/K16apache /etc/rc1.d/k16apache

# mv /etc/rc2.d/K16apache /etc/rc2.d/k16apache

# mv /etc/rc3.d/S50apache /etc/rc3.d/s50apache

# mv /etc/rcS.d/K16apache /etc/rcS.d/k16apache

# ls -l /etc/rc?.d/*apache
/etc/rc0.d/k16apache/etc/rc1.d/k16apache/etc/rc2.d/k16apache/etc/rc3.d/s50apache/etc/rcS.d/k16ap
```

Man Pages

New man pages are included for each data service supplied with Sun Cluster 3.0 software. The data service man pages include: `SUNW.apache(5)`, `SUNW.dns(5)`, `SUNW.iws(5)`, `SUNW.nfs(5)`, `SUNW.nslldap(5)`, `SUNW.oracle_listener(5)`, `SUNW.oracle_server(5)`, `SUNW.HAStorage(5)` and `scalable_service(5)`. These man pages describe the standard and extension properties that these data services use.

Known Problems With the Sun Management Center GUI

This section describes known problems with the Sun Cluster 3.0 module of the Sun Management Center GUI.

Certain Types of Ultra Servers are Not Recognized by Sun Management Center

Symptoms

- A blue or yellow folder icon appears in the Sun Management Center window to represent the cluster node under its domain instead of the expected machine icon.
- The Load Modules menu choice is grayed out when the node is shown in the Details Window.

Confirmation of Problem/Start of Workaround

1. Close the Details Window.
2. From the Sun Management Center Window, choose File->Console Messages.
3. Double-click the folder icon representing the unrecognized cluster node.
4. Look in the console messages window for a line reading `...family definition file missing for...`

Workaround

1. On the Sun Management Center server, change to the directory holding family files.

```
# cd /opt/SUNWsymon/classes/base/console/cfg
```

2. Create a symbolic link to the closest available family-j.x file.

For example, if the missing file line read `...missing for sun4u-Sun-Ultra-450-family-j.x...`, create a link from `sun4u-Sun-Enterprise-450-family-j.x` to `sun4u-Sun-Ultra-450-family-j.x`.

```
# ln -s sun4u-Sun-Enterprise-450-family-j.x sun4u-Sun-Ultra-450-family-j.x
```

3. Exit the console, and restart it.

Alternate Method for Determining Names to Symbolic Link

1. Double-click the unrecognized cluster node to bring up its Details Window.

2. Click the Info tab.

3. Search for the Entity Family entry in the Properties table.

The value will probably be truncated, so let the mouse pointer linger over the value field. The complete name (for example, sun4u-Sun-Ultra-450) appears in the tooltip.

4. Append `-family-j.x` to determine the link name to create.

Sun Cluster Installation Configuration Worksheets and Examples

This appendix provides worksheets for planning various components of your cluster configuration. Examples of completed worksheets are also provided for your reference. You will enter information from your completed worksheets during installation and configuration of your Sun Cluster configuration. *Sun Cluster 3.0 Installation Guide* contains guidelines to follow when completing these worksheets.

This appendix contains the following worksheet information:

- “Installation Configuration Worksheets” on page 43
- “Installation Configuration Worksheet Examples” on page 53

Appendix B contains configuration worksheets for resources, resource types, and resource groups.

Installation Configuration Worksheets

This section provides the following planning worksheets:

- “Cluster and Node Names Worksheet” on page 45
- “Cluster Interconnect Worksheet” on page 46
- “Public Networks Worksheet” on page 47
- “Local Devices Worksheet” on page 48
- “Local File System Layout Worksheet” on page 49
- “Disk Device Group Configurations Worksheet” on page 50
- “Volume Manager Configurations Worksheet” on page 51

- “Metadevices Worksheet (Solstice DiskSuite)” on page 52

You might need to make multiple copies of a worksheet to accommodate all the components in your cluster configuration.

Cluster and Node Names Worksheet

Cluster name _____

Privatenetworkaddress: _____ . 0 . 0 (default: 172.16.0.0)

Privatenetworkmask: 255.255.____.____ (default: 255.255.0.0)

Sponsor node name _____

Privatehostname

Default name: clusternode____-priv
node ID

Change to: _____
(optional)

Additional nodes

<p>Node name _____</p> <p>Privatehostname</p> <p>Default name: clusternode____-priv node ID</p> <p>Change to: _____ (optional)</p>	<p>Node name _____</p> <p>Privatehostname</p> <p>Default name: clusternode____-priv node ID</p> <p>Change to: _____ (optional)</p>
<p>Node name _____</p> <p>Privatehostname</p> <p>Default name: clusternode____-priv node ID</p> <p>Change to: _____ (optional)</p>	<p>Node name _____</p> <p>Privatehostname</p> <p>Default name: clusternode____-priv node ID</p> <p>Change to: _____ (optional)</p>
<p>Node name _____</p> <p>Privatehostname</p> <p>Default name: clusternode____-priv node ID</p> <p>Change to: _____ (optional)</p>	<p>Node name _____</p> <p>Privatehostname</p> <p>Default name: clusternode____-priv node ID</p> <p>Change to: _____ (optional)</p>
<p>Node name _____</p> <p>Privatehostname</p> <p>Default name: clusternode____-priv node ID</p> <p>Change to: _____ (optional)</p>	Empty space for the fourth node's second column

Cluster Interconnect Worksheet

Adapters

Cabling

Junctions

Draw lines between cable endpoints

Node name _____

Adapter Name	Transport Type

Junction name _____

Junction type _____

Port Name

Node name _____

Adapter Name	Transport Type

Junction name _____

Junction type _____

Port Name

Node name _____

Adapter Name	Transport Type

Node name _____

Adapter Name	Transport Type

Public Networks Worksheet

Node name _____

Primary hostname _____

NAFO group number: *nafo*____

Adapter name _____

Backup adapter(s): _____
(*optional*)

Network name _____

Secondary hostname _____

NAFO group number: *nafo*____

Adapter name _____

Backup adapter(s): _____
(*optional*)

Network name _____

Secondary hostname _____

NAFO group number: *nafo*____

Adapter name _____

Backup adapter(s): _____
(*optional*)

Network name _____

Secondary hostname _____

NAFO group number: *nafo*____

Adapter name _____

Backup adapter(s): _____
(*optional*)

Network name _____

Node name _____

Primary hostname _____

NAFO group number: *nafo*____

Adapter name _____

Backup adapter(s): _____
(*optional*)

Network name _____

Secondary hostname _____

NAFO group number: *nafo*____

Adapter name _____

Backup adapter(s): _____
(*optional*)

Network name _____

Secondary hostname _____

NAFO group number: *nafo*____

Adapter name _____

Backup adapter(s): _____
(*optional*)

Network name _____

Secondary hostname _____

NAFO group number: *nafo*____

Adapter name _____

Backup adapter(s): _____
(*optional*)

Network name _____

Local Devices Worksheet

Node name _____

Local disks

Disk name _____ Size _____

Disk name _____ Size _____

Disk name _____ Size _____

Disk name _____ Size _____

Disk name _____ Size _____

Disk name _____ Size _____

Disk name _____ Size _____

Disk name _____ Size _____

Other local devices

Devicetype _____ Name _____

Devicetype _____ Name _____

Devicetype _____ Name _____

Devicetype _____ Name _____

Node name _____

Local disks

Disk name _____ Size _____

Disk name _____ Size _____

Disk name _____ Size _____

Disk name _____ Size _____

Disk name _____ Size _____

Disk name _____ Size _____

Disk name _____ Size _____

Disk name _____ Size _____

Other local devices

Devicetype _____ Name _____

Devicetype _____ Name _____

Devicetype _____ Name _____

Devicetype _____ Name _____

Local File System Layout Worksheet

Node name _____

Mirrored root

Volume Name	Component	Component	File System	Size
			/	
			swap	
			/globaldevices	

Non-mirrored root

Device Name	File System	Size
	/	
	swap	
	/globaldevices	

Disk Device Group Configurations Worksheet

Volume manager: _____

Disk group/diskset name _____

Node names (1) _____ (2) _____ (3) _____ (4) _____

(5) _____ (6) _____ (7) _____ (8) _____

Ordered priority? Yes No

Failback? Yes No

Disk group/diskset name _____

Node names (1) _____ (2) _____ (3) _____ (4) _____

(5) _____ (6) _____ (7) _____ (8) _____

Ordered priority? Yes No

Failback? Yes No

Disk group/diskset name _____

Node names (1) _____ (2) _____ (3) _____ (4) _____

(5) _____ (6) _____ (7) _____ (8) _____

Ordered priority? Yes No

Failback? Yes No

Volume Manager Configurations Worksheet

Volume manager: _____

Name	Type	Component	Component

Installation Configuration Worksheet Examples

This section provides the following completed examples of the configuration planning worksheets in “Installation Configuration Worksheets” on page 43:

Note - The data used in these examples is intended as a guideline only and does not represent a complete configuration of a functional cluster.

- “Example: Cluster and Node Names” on page 54
- “Example: Cluster Interconnect” on page 55
- “Example: Public Networks” on page 56
- “Example: Local Devices” on page 57
- “Example: Local File System Layout—With Mirrored Root” on page 58
- “Example: Local File System Layout—Without Mirrored Root” on page 59
- “Example: Disk Device Group Configurations” on page 60
- “Example: Volume Manager Configurations” on page 61
- “Example: Metadevices (Solstice DiskSuite)” on page 62

Example: Cluster and Node Names

Cluster name scluster

Private network address: 172 . 16 . 0 . 0 (default: 172.16.0.0)

Private network mask: 255 . 255 . 0 . 0 (default: 255.255.0.0)

Sponsor node name phys-schost-1

Private hostname

Default name: clusternode 1 -priv
node ID

Change to: phys-schost-1-priv
(optional)

Additional nodes

<p>Node name <u> phys-schost-2 </u></p> <p>Private hostname</p> <p>Default name: clusternode<u> 2 </u>-priv <i>node ID</i></p> <p>Change to: <u> phys-schost-2-priv </u> <i>(optional)</i></p>	<p>Node name _____</p> <p>Private hostname</p> <p>Default name: clusternode_____priv <i>node ID</i></p> <p>Change to: _____ <i>(optional)</i></p>
<p>Node name _____</p> <p>Private hostname</p> <p>Default name: clusternode_____priv <i>node ID</i></p> <p>Change to: _____ <i>(optional)</i></p>	<p>Node name _____</p> <p>Private hostname</p> <p>Default name: clusternode_____priv <i>node ID</i></p> <p>Change to: _____ <i>(optional)</i></p>
<p>Node name _____</p> <p>Private hostname</p> <p>Default name: clusternode_____priv <i>node ID</i></p> <p>Change to: _____ <i>(optional)</i></p>	<p>Node name _____</p> <p>Private hostname</p> <p>Default name: clusternode_____priv <i>node ID</i></p> <p>Change to: _____ <i>(optional)</i></p>
<p>Node name _____</p> <p>Private hostname</p> <p>Default name: clusternode_____priv <i>node ID</i></p> <p>Change to: _____ <i>(optional)</i></p>	

Example: Cluster Interconnect

Adapters

Cabling

Junctions

Draw lines between cable endpoints

Node name phys-schost1

Adapter Name	Transport Type
hme0	dlpi
hme1	dlpi

Junction name switch1

Junction type switch

PortName
1
2

Node name phys-schost2

Adapter Name	Transport Type
hme0	dlpi
hme1	dlpi

Junction name switch2

Junction type switch

PortName
1
2

Node name _____

Adapter Name	Transport Type

Node name _____

Adapter Name	Transport Type

Example: Public Networks

Node name phys-schost-1

Primary hostname phys-schost-1

NAFO group number: nafo_0

Adapter name qfe0

Backup adapter(s): qfe4
(optional)

Network name net-85

Secondary hostname phys-schost-1-86

NAFO group number: nafo_1

Adapter name qfe1

Backup adapter(s): qfe5
(optional)

Network name net-86

Secondary hostname _____

NAFO group number: nafo_

Adapter name _____

Backup adapter(s): _____
(optional)

Network name _____

Secondary hostname _____

NAFO group number: nafo_

Adapter name _____

Backup adapter(s): _____
(optional)

Network name _____

Node name phys-schost-2

Primary hostname phys-schost-2

NAFO group number: nafo_0

Adapter name qfe0

Backup adapter(s): qfe4
(optional)

Network name net-85

Secondary hostname phys-schost-2-86

NAFO group number: nafo_1

Adapter name qfe1

Backup adapter(s): qfe5
(optional)

Network name net-86

Secondary hostname _____

NAFO group number: nafo_

Adapter name _____

Backup adapter(s): _____
(optional)

Network name _____

Secondary hostname _____

NAFO group number: nafo_

Adapter name _____

Backup adapter(s): _____
(optional)

Network name _____

Example: Local Devices

Node name phys-schost-1

Local disks

Disk name c0t0d0 Size 2G

Disk name c0t1d0 Size 2G

Disk name c1t0d0 Size 2G

Disk name c1t1d0 Size 2G

Disk name _____ Size _____

Disk name _____ Size _____

Disk name _____ Size _____

Disk name _____ Size _____

Other local devices

Devicetype tape Name /dev/rmt0

Devicetype _____ Name _____

Devicetype _____ Name _____

Devicetype _____ Name _____

Node name _____

Local disks

Disk name _____ Size _____

Disk name _____ Size _____

Disk name _____ Size _____

Disk name _____ Size _____

Disk name _____ Size _____

Disk name _____ Size _____

Disk name _____ Size _____

Disk name _____ Size _____

Other local devices

Device type _____ Name _____

Device type _____ Name _____

Device type _____ Name _____

Device type _____ Name _____

Example: Local File System Layout—With Mirrored Root

Node name phys-schost-1

Mirrored root

Volume Name	Component	Component	File System	Size
d1	c0t0d0s0	c1t0d0s0	/	1168MB
d2	c0t0d0s1	c1t0d0s1	swap	750MB
d3	c0t0d0s3	c1t0d0s3	/globaldevices	100MB
d7	c0t0d0s7	c1t0d0s7	SDS replica	10MB

Non-mirrored root

Device Name	File System	Size
	/	
	swap	
	/globaldevices	

Example: Local File System Layout—Without Mirrored Root

Node name phys-schost-1

Mirrored root

Volume Name	Component	Component	File System	Size
			/	
			swap	
			/globaldevices	

Non-mirrored root

Device Name	File System	Size
c0t0d0s0	/	1168 MB
c0t0d0s1	swap	750MB
c0t0d0s3	/globaldevices	100MB
c0t0d0s7	SDS replica	10MB

Example: Disk Device Group Configurations

Volume manager: Solstice DiskSuite

Disk group/diskset name dg-schost-1

Node names (1) phys-schost-1 (2) phys-schost-2 (3) _____ (4) _____

(5) _____ (6) _____ (7) _____ (8) _____

Ordered priority? Yes No

Failback? Yes No

Disk group/diskset name _____

Node names (1) _____ (2) _____ (3) _____ (4) _____

(5) _____ (6) _____ (7) _____ (8) _____

Ordered priority? Yes No

Failback? Yes No

Disk group/diskset name _____

Node names (1) _____ (2) _____ (3) _____ (4) _____

(5) _____ (6) _____ (7) _____ (8) _____

Ordered priority? Yes No

Failback? Yes No

Example: Volume Manager Configurations

Volume manager: Solstice DiskSuite

Name	Type	Component	Component
dg-schost-1/d0	trans	dg-schost-1/d1	dg-schost-1/d4
dg-schost-1/d1	mirror	c0t0d0s4	c4t4d0s4
dg-schost-1/d4	mirror	c0t0d2s5	c4t4d2s5

Data Service Configuration Worksheets and Examples

This appendix provides worksheets for planning resource-related components of your cluster configuration. Examples of completed worksheets are also provided for your reference.

This appendix contains the following worksheet information:

- “Configuration Worksheets” on page 63
- “Configuration Worksheet Examples” on page 69

Configuration Worksheets

This section provides worksheets for planning your:

- “Resource Types” on page 64
- “Resource Groups—Failover” on page 65
- “Resource Groups—Scalable” on page 66
- “Network Resources” on page 67

You might need to make multiple copies of a worksheet to accommodate all the components in your cluster configuration.

Resource Types

(other than logical host or shared address)
Indicate the nodes on which the resource type will run.

Resource type name _____

Node names _____

Resource type name _____

Node names _____

Resource type name _____

Node names _____

Resource type name _____

Node names _____

Resource type name _____

Node names _____

Resource type name _____

Node names _____

Resource Groups—Failover

Resource group name _____

(Must be unique within the cluster.)

Function of this resource group _____

Failback? Yes No

(Will this resource group switch back to the primary node, after the primary node has failed and been restored?)

Node names (1) _____ (2) _____ (3) _____ (4) _____
(ordered list)
(5) _____ (6) _____ (7) _____ (8) _____

(Indicate the cluster nodes that may host this resource group. The first node in this list should be the primary, with others being the secondaries. The order of the secondaries will indicate preference for becoming primaries.)

Disk device groups upon which this resource group depends _____

(If the resources in this resource group need to create files for administrative purposes, include the subdirectory they should use.)

Resource Groups—Scalable

Resource group name _____

(Must be unique within the cluster.)

Function of this resource group _____

Maximum number of primaries _____

Desired number of primaries _____

Failback? Yes No

(Will this resource group switch back to the primary node, after the primary node has failed)

Node names (1) _____ (2) _____ (3) _____ (4) _____

(5) _____ (6) _____ (7) _____ (8) _____

(Indicate the cluster nodes that may host this resource group. The first node in this list should be the primary, with others being the secondaries. The order of the secondaries will indicate preference for becoming primaries.)

Dependencies _____

(Does this resource depend upon another resource group.)

Configuration Worksheet Examples

This section provides the following completed examples of the configuration planning worksheets in “Configuration Worksheets” on page 63:

Note - The data used in these examples are intended as guidelines only and do not represent a complete configuration of a functional cluster.

- “Example: Resource Types” on page 70
- “Example: Resource Groups—Scalable” on page 71
- “Example: Resource Groups—Failover” on page 72
- “Example: Network Resources—LogicalHostname” on page 73
- “Example: Resources—SharedAddress” on page 74
- “Example: Resources—Application” on page 75

Example: Resource Types

(other than logical host or shared address)

Resource type name SUNW.nshttp

Node names phys-schost-1 phys-schost-2 _____

Resource type name SUNW.oracle listener

Node names phys-schost-1 phys-schost-2 _____

Resource type name SUNW.oracle_server

Node names phys-schost-1 phys-schost-2 _____

Resource type name _____

Node names _____

Resource type name _____

Node names _____

Resource type name _____

Node names _____

Example: Resource Groups—Scalable

Resource group name rg-http

(Must be unique within the cluster.)

Function of this resource group Contains the web server resources

Maximum number of primaries 2

Desired number of primaries 2

Failback? Yes No

(Will this resource group switch back to the primary node, after the primary node has failed)

Node names (1) phys-schost-1 (2) phys-schost-2 (3) _____ (4) _____
(ordered list)

(5) _____ (6) _____ (7) _____ (8) _____

(Indicate the cluster nodes that may host this resource group. The first node in this list should be the primary, with others being the secondaries. The order of the secondaries will indicate preference for becoming primaries.)

Dependencies rg-shared

(Does this resource depend upon another resource group.)

Resource group name rg-shared

(Must be unique within the cluster.)

Function of this resource group Contains the shared address resources

Maximum number of primaries 1

Desired number of primaries 1

Failback? Yes No

(Will this resource group switch back to the primary node, after the primary node has failed)

Node names (1) phys-schost-1 (2) phys-schost-2 (3) _____ (4) _____
(ordered list)

(5) _____ (6) _____ (7) _____ (8) _____

(Indicate the cluster nodes that may host this resource group. The first node in this list should be the primary, with others being the secondaries. The order of the secondaries will indicate preference for becoming primaries.)

Dependencies _____

Example: Resource Groups—Failover

Resource group name rg-oracle

(Must be unique within the cluster.)

Function of this resource group Contains the Oracle resources

Failback? Yes No

(Will this resource group switch back to the primary node, after the primary node has failed and been restored?)

Node names (1) phys-schost-1 (2) phys-schost-2 (3) _____ (4) _____
(ordered list)
(5) _____ (6) _____ (7) _____ (8) _____

(Indicate the cluster nodes that may host this resource group. The first node in this list should be the primary, with others being the secondaries. The order of the secondaries will indicate preference for becoming primaries.)

Disk device groups upon which this resource group depends schost1-dg

(If the resources in this resource group need to create files for administrative purposes, include the subdirectory they should use.)

