



Sun Cluster 3.2 Release Notes for Solaris OS



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

Part No: 819-6611-10
December 2006, Revision A

Copyright 2006 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more U.S. patents or pending patent applications in the U.S. and in other countries.

U.S. Government Rights – Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

This distribution may include materials developed by third parties.

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, the Solaris logo, the Java Coffee Cup logo, docs.sun.com, Java, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon 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.

Products covered by and information contained in this publication are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical or biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

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 2006 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Sun Microsystems, Inc. détient les droits de propriété intellectuelle relatifs à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et ce sans limitation, ces droits de propriété intellectuelle peuvent inclure un ou plusieurs brevets américains ou des applications de brevet en attente aux États-Unis et dans d'autres pays.

Cette distribution peut comprendre des composants développés par des tierces personnes.

Certains composants de ce produit peuvent être dérivées du logiciel Berkeley BSD, licenciés par l'Université de Californie. UNIX est une marque déposée aux États-Unis et dans d'autres pays; elle est licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, le logo Solaris, le logo Java Coffee Cup, docs.sun.com, Java et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux États-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 États-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.

Les produits qui font l'objet de cette publication et les informations qu'il contient sont régis par la législation américaine en matière de contrôle des exportations et peuvent être soumis au droit d'autres pays dans le domaine des exportations et importations. Les utilisations finales, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes chimiques ou biologiques ou pour le nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou réexportations vers des pays sous embargo des États-Unis, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exclusive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

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

Preface	5
Sun Cluster 3.2 Release Notes for Solaris OS	9
What's New in the Sun Cluster 3.2 Software	9
New Features and Functionality	9
Compatibility Issues	15
Features Nearing End of Life	16
Solaris 10 11/06 Operating System	16
Solaris Volume Manager GUI	17
Loopback File System (LOFS)	17
HA-Apache Configuration Wizard	17
Accessibility Features for People With Disabilities	17
Commands Modified in This Release	18
Object-Oriented Command Line Interface	18
scinstall Command	18
scconf Command	18
Resource Properties	19
Product Name Changes	19
Supported Products	19
Sun Cluster Security Hardening	24
Known Issues and Bugs	25
Administration	26
Data Services	29
Installation	31
Localization	33
Quorum Server	35
Runtime	35
Upgrade	37

Patches and Required Firmware Levels	38
Patch Management Tools	39
SunSolve Online	39
Sun Cluster 3.2 Documentation	40
Searching Sun Product Documentation	40
Sun Cluster 3.2 Software Manuals for Solaris OS	41
Sun Cluster 3.2 Data Service Manuals for Solaris OS (SPARC Platform Edition)	41
Sun Cluster 3.2 Data Service Manuals for Solaris OS (x86 Platform Edition)	43
Sun Cluster 3.1 – 3.2 Hardware Collection for Solaris OS (SPARC Platform Edition)	44
Sun Cluster 3.1 — 3.2 Hardware Collection for Solaris OS (x86 Platform Edition)	45
Documentation Issues	45
Software Installation Guide	45
▼ How to Recover From a Partially Completed Dual-Partition Upgrade on a SPARC Based System	47
▼ How to Recover From a Partially Completed Dual-Partition Upgrade on an x86 Based System	48
Sun Cluster Data Services Planning and Administration Guide	50
System Administration Guide	50
Data Services Developer's Guide	52
Data Service for SAP Guide	53
Quorum Server User's Guide	56
Man Pages	56

Preface

Sun Cluster 3.2 Release Notes for Solaris OS This document lists product features, requirements, and restrictions for Sun™ Cluster 3.2 software. It also describes open bugs and other known problems.

Note – In this document, the term “x86” refers to the Intel 32-bit family of microprocessor chips and compatible microprocessor chips made by AMD.

This document is intended for experienced system administrators with extensive knowledge of Sun software and hardware. This document is not to be used as a planning or presales guide.

The instructions in this book assume knowledge of the Solaris™ operating system and expertise with the volume manager software used with Sun Cluster software.

Note – Sun Cluster software runs on two platforms, SPARC® and x86. The information in this document pertains to both platforms unless otherwise specified in a special chapter, section, note, bulleted item, figure, table, or example.

Using UNIX Commands

This document contains information on commands specific to administering a Sun Cluster configuration. This document might not contain complete information on basic UNIX® commands and procedures.

See one or more of the following for this information:

- Online documentation for the Solaris software
- Other software documentation that you received with your system
- Solaris operating system man pages

Typographic Conventions

The following table describes the typographic conventions that are used in this book.

TABLE P-1 Typographic Conventions

Typeface	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name%</code> you have mail.
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name%</code> su Password:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <i>rm filename</i> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . <i>A cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

The following table shows the default UNIX system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

TABLE P-2 Shell Prompts

Shell	Prompt
C shell	<code>machine_name%</code>
C shell for superuser	<code>machine_name#</code>
Bourne shell and Korn shell	<code>\$</code>
Bourne shell and Korn shell for superuser	<code>#</code>

Related Documentation

Information about related Sun Cluster topics is available in the documentation that is listed in the following table. All Sun Cluster documentation is available at <http://docs.sun.com>.

Topic	Documentation
Overview	<i>Sun Cluster Overview for Solaris OS</i>
Concepts	<i>Sun Cluster Concepts Guide for Solaris OS</i>
Hardware installation and administration	<i>Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS</i> Individual hardware administration guides
Software installation	<i>Sun Cluster Software Installation Guide for Solaris OS</i>
Data service installation and administration	<i>Sun Cluster Data Services Planning and Administration Guide for Solaris OS</i> Individual data service guides
Data service development	<i>Sun Cluster Data Services Developer's Guide for Solaris OS</i>
System administration	<i>Sun Cluster System Administration Guide for Solaris OS</i>
Error messages	<i>Sun Cluster Error Messages Guide for Solaris OS</i>
Command and function references	<i>Sun Cluster Reference Manual for Solaris OS</i>

Documentation, Support, and Training

The Sun web site provides information about the following additional resources:

- [Documentation \(http://www.sun.com/documentation/\)](http://www.sun.com/documentation/)
- [Support \(http://www.sun.com/support/\)](http://www.sun.com/support/)
- [Training \(http://www.sun.com/training/\)](http://www.sun.com/training/)

Getting Help

Contact your service provider if you have problems installing or using Sun Cluster. Provide the following information to your service provider.

- Your name and email address
- Your company name, address, and phone number
- The model and serial numbers of your systems
- The release number of the operating environment, for example Solaris 9
- The release number of Sun Cluster software, for example, Sun Cluster 3.2

Use the following commands to gather information on your system for your service provider:

Command	Function
<code>prtconf -v</code>	Displays the size of the system memory and reports information about peripheral devices
<code>psrinfo -v</code>	Displays information about processors
<code>showrev -p</code>	Reports which patches are installed
<code>SPARC: prtdiag -v</code>	Displays system diagnostic information
<code>/usr/cluster/bin/scinstall -pv</code>	Displays Sun Cluster release and package version information

Also, have available the contents of the `/var/adm/messages` file.

Sun Cluster 3.2 Release Notes for Solaris OS

This document provides the following information for Sun™ Cluster 3.2 software.

- “What's New in the Sun Cluster 3.2 Software” on page 9
- “Compatibility Issues” on page 15
- “Commands Modified in This Release” on page 18
- “Product Name Changes” on page 19
- “Supported Products” on page 19
- “Known Issues and Bugs” on page 25
- “Patches and Required Firmware Levels” on page 38
- “Sun Cluster 3.2 Documentation” on page 40
- “Documentation Issues” on page 45

What's New in the Sun Cluster 3.2 Software

This section provides information related to new features, functionality, and supported products in the Sun Cluster 3.2 software. This section also provides information on any restrictions that are introduced in this release.

New Features and Functionality

This section describes each of the following new features provided in the Sun Cluster 3.2 software.

- “New Sun Cluster Object-Oriented Command Set” on page 10
- “Oracle RAC 10g Improved Integration and Administration” on page 10
- “Data Service Configuration Wizards” on page 10
- “Flexible IP Address Scheme” on page 11
- “Sun Cluster Support for Service Management Facility Services” on page 11
- “Extended Flexibility for Fencing Protocol” on page 11
- “Sun Cluster Quorum Server” on page 12

- “Disk-Path Failure Handling” on page 12
- “HASStoragePlus Availability Improvements” on page 12
- “Solaris Zones Expanded Support” on page 12
- “ZFS” on page 13
- “HDS TrueCopy Campus Cluster” on page 13
- “Multi-Terabyte Disk and Extensible Firmware Interface (EFI) Label Support” on page 13
- “Extended Support for VERITAS Software Components” on page 13
- “Quota Support” on page 14
- “Oracle DataGuard Support” on page 14
- “Dual-Partition Upgrade” on page 14
- “Live Upgrade” on page 14
- “Optional Sun Cluster Manager Installation” on page 14
- “SNMP Event MIB” on page 14
- “Command Logging” on page 15
- “Workload System-Resource Monitoring” on page 15
- “Automatic Creation of Multiple-Adapter IPMP Groups by `scinstall`” on page 15

New Sun Cluster Object-Oriented Command Set

The new Sun Cluster command-line interface includes a separate command for each cluster object type and uses consistent subcommand names and option letters. The new Sun Cluster command set also supports short and long command names. The command output provides improved help and error messages as well as more readable status and configuration reports. In addition, some commands include export and import options with the use of portable XML-based configuration files. These options allow you to replicate a portion of, or the entire, cluster configuration, which speeds up partial or full configuration cloning. See the `Intro(1CL)` man page for more information.

Oracle RAC 10g Improved Integration and Administration

Sun Cluster Oracle RAC package installation, as well as configuration, is now integrated in the Sun Cluster procedures. New Oracle RAC-specific resource types and properties can be used for finer-grained control.

Oracle RAC extended manageability, which is provided by the `ScalDeviceGroup` and `ScalMountPoint` resource types, leads to easier set up of Oracle RAC within Sun Cluster configurations, as well as improved diagnosability and availability. See *Sun Cluster Data Service for Oracle RAC Guide for Solaris OS* for more information.

Data Service Configuration Wizards

Sun Cluster provides new data service configuration wizards that simplify configuration of popular applications through automatic discovery of parameter choices and immediate validation. The Sun Cluster data service configuration wizards are provided in the following two formats:

- Sun Cluster Manager GUI
- `clsetup` command-line interface

The following data services are supported in the Sun Cluster Manager GUI format:

- HA-Oracle
- HA-NFS
- HA-Apache
- HA-SAP

The `clsetup` command-line interface format supports all applications that are supported by Sun Cluster Manager plus Oracle RAC .

See the Sun Cluster documentation for each of the supported data services for more information.

Flexible IP Address Scheme

Sun Cluster software now allows a reduced range of IP addresses for its private interconnect. In addition, you can now customize the IP base address and its range during or after installation.

These changes to the IP address scheme facilitate integration of Sun Cluster environments in existing networks with limited or regulated address spaces. See “How to Change the Private Network Address or Address Range of an Existing Cluster” in *Sun Cluster System Administration Guide for Solaris OS* in the *Sun Cluster System Administration Guide for Solaris OS* for more information.

Sun Cluster Support for Service Management Facility Services

Sun Cluster software now integrates tightly with Solaris 10 OS Service Management Facility (SMF) and enables the encapsulation of SMF-controlled applications in the Sun Cluster resource management model. Local service-level life-cycle management continues to be operated by SMF, while whole resource level cluster-wide failure-handling operations (node, storage) are carried out by Sun Cluster software.

Moving applications from a single-node Solaris 10 OS environment to a multi-node Sun Cluster environment enables increased availability while requiring minimal effort. See “Enabling Solaris SMF Services to Run With Sun Cluster” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* in the *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* for more information.

Extended Flexibility for Fencing Protocol

This new functionality allows the customization of the default fencing protocol. Choices include SCSI-3, SCSI-2, or per-device discovery.

This flexibility enables the default usage of SCSI-3, a more recent protocol, for better support for multipathing, easier integration with non-Sun storage, and shorter recovery times on newer storage while still supporting the Sun Cluster 3.0 or 3.1 behavior and SCSI-2 for older devices. See “Administering the SCSI Protocol Settings for Storage Devices” in *Sun Cluster System Administration Guide for Solaris OS* for more information.

Sun Cluster Quorum Server

A new quorum device option is now available in Sun Cluster software. Instead of using a shared disk and SCSI reservation protocols, it is now possible to use a Solaris server outside of the cluster to run a quorum-server module, which supports an atomic reservation protocol over TCP/IP. This support enables faster failover time and also lowers deployment costs: it removes the need of a shared quorum disk for any scenario where quorum is required (two-node) or desired. See *Sun Cluster Quorum Server User's Guide* for more information.

Disk-Path Failure Handling

Sun Cluster software can now be configured to automatically reboot a node if all its path to shared disks have failed. Faster reaction in case of severe disk-path failure enables improved availability. See “Administering Disk-Path Monitoring” in *Sun Cluster System Administration Guide for Solaris OS* in the *Sun Cluster System Administration Guide for Solaris OS* for more information.

HAStoragePlus Availability Improvements

HAStoragePlus mount points are now created automatically in case of mount failure. This feature eliminates failure-to-fail over cases, thus improving availability of the environment.

Solaris Zones Expanded Support

Sun Cluster software now supports the following data services in Solaris non-global zones.

- Sun Cluster Data Service for Apache
- Sun Cluster Data Service for Apache Tomcat
- Sun Cluster Data Service for DHCP
- Sun Cluster Data Service for Domain Name Service (DNS)
- Sun Cluster Data Service for Kerberos
- Sun Cluster Data Service for MySQL
- Sun Cluster Data Service for N1 Grid Service Provisioning Server
- Sun Cluster Data Service for Oracle
- Sun Cluster Data Service for Oracle Application Server
- Sun Cluster Data Service for PostgreSQL
- Sun Cluster Data Service for Samba
- Sun Cluster Data Service for Sun Java System Application Server
- Sun Cluster Data Service for Sun Java System Message Queue Server
- Sun Cluster Data Service for Sun Java System Web Server

This allows the combination of the benefits of application containment that is offered by Solaris zones and the increased availability that is provided by Sun Cluster software. See the Sun Cluster documentation for the appropriate data services for more information.

ZFS

ZFS is supported as a highly available local file system in Sun Cluster 3.2. ZFS with Sun Cluster software offers a best-class file system solution combining high availability, data integrity, performance, and scalability, covering the needs of the most demanding environments.

Continuous enhancements are being added to ZFS for optimizing performance with all workloads, especially database transactions. Ensure that you have the latest ZFS patches installed and that your configuration is optimized for your specific type of workload.

HDS TrueCopy Campus Cluster

Sun Cluster-based campus clusters now support HDS TrueCopy controller-based replication, allowing for automated management of TrueCopy configurations. Sun Cluster software handles automatically and transparently the switch to the secondary campus site in case of failover, making this procedure less error-prone and improving the overall availability of the solution. This new remote data-replication infrastructure allows Sun Cluster software to support new configurations for customers who have been standardizing on specific replication infrastructure like TrueCopy, and for places where host-based replication is not a viable solution because of distance or application incompatibility.

This new combination brings improved availability and less complexity while lowering cost. Sun Cluster software can make use of existing TrueCopy customer replication infrastructure, limiting the need for additional replication solutions.

Multi-Terabyte Disk and Extensible Firmware Interface (EFI) Label Support

Sun Cluster configurations now support disks with a capacity over 1TB which use a new Extensible Firmware Interface (EFI) disk format. This format is required for multi-terabyte disks but can also be used with smaller capacity disks. This extends the supported Sun Cluster configurations to environments with high-end storage requirements.

Extended Support for VERITAS Software Components

VERITAS Volume Manager and File System, part of VERITAS Storage Foundation 5.0, are now supported on SPARC platforms as well as VERITAS Volume Manager 4.1 with Solaris 10 OS on x86/x64 platforms.

VERITAS Volume Replicator (VVR) 5.0 and VERITAS Fast Mirror Resynchronization (FMR) 4.1 and 5.0, part of VERITAS FlashSnap, can now be used in Sun Cluster environments on SPARC platforms.

Quota Support

Quota management can now be used with HAStoragePlus on local UFS file systems for better control of resource consumption.

Oracle DataGuard Support

Customers are now able to operate Oracle DataGuard data replication configurations under Sun Cluster control. Sun Cluster software now offers improved usability for Oracle deployments including DataGuard data replication software. For more information, see *Sun Cluster Data Service for Oracle RAC Guide for Solaris OS*.

Dual-Partition Upgrade

With this new software swap feature the upgrade process is greatly simplified. Any component of the software stack along with Sun Cluster software can be upgraded in one step: Solaris operating system, Sun Cluster software, file systems, volume managers, applications, and data services. This automation lowers the risk of human errors during cluster upgrade and minimizes the service outage that occur for a standard cluster upgrade.

Live Upgrade

The Live Upgrade method can now be used with Sun Cluster software. This method reduces system downtime of a node during upgrade as well as unnecessary reboots, therefore lowering the required maintenance window where the service is at risk.

At the time of publication, Live Upgrade can be used only if your Sun Cluster installation uses Solaris Volume Manager for managing the storage or disk groups. Live Upgrade does not currently support VxVM. See “[Upgrade](#)” on [page 37](#) for more information.

Any Live Upgrade from Solaris 8 to Solaris 9 requires SVM patch 116669-18 to be applied before rebooting from the alternate root.

Optional Sun Cluster Manager Installation

Installation of Sun Cluster Manager, the Sun Cluster management GUI, is now optional. This change removes web-based access to the cluster, to comply with potential security rules. See “How to Install Sun Cluster Framework and Data-Service Software Packages” in *Sun Cluster Software Installation Guide for Solaris OS* in the *Sun Cluster Software Installation Guide for Solaris OS* for information about deselecting Sun Cluster Manager at installation time.

SNMP Event MIB

Sun Cluster software includes a new Sun Cluster SNMP event mechanism as well as a new SNMP MIB. They now allow third-party SNMP management applications to directly register with Sun Cluster software and receive timely notifications of cluster events. Fine-grained event

notification and direct integration with third-party enterprise-management framework through standard SNMP support allow proactive monitoring and increase availability. See “Creating, Setting Up, and Managing the Sun Cluster SNMP Event MIB” in *Sun Cluster System Administration Guide for Solaris OS* for more information.

Command Logging

Command information can now be logged within Sun Cluster software. This ability facilitates diagnostics of cluster failures and provides history of the administration actions for archiving or replication. For more information, see “How to View the Contents of Sun Cluster Command Logs” in *Sun Cluster System Administration Guide for Solaris OS* in the *Sun Cluster System Administration Guide for Solaris OS*.

Workload System-Resource Monitoring

Sun Cluster software offers new system-resources utilization measurement and visualization tools, including fine-grained measurement of consumptions per node, resource, and resource group. These new tools provide historical data as well as threshold management and CPU reservation and control. This improved control allows for better management of service level and capacity.

Automatic Creation of Multiple-Adapter IPMP Groups by `scinstall`

The interactive `scinstall` utility no longer automatically creates a single-adapter IPMP group on each public-network adapter in the cluster. Instead, the utility configures a multiple-adapter IPMP group for each set of public-network adapters that uses the same subnet, except for adapters that are already configured in an IPMP group. For more information about this and other changes to IPMP group policies, see “Public Networks” in *Sun Cluster Software Installation Guide for Solaris OS*.

Compatibility Issues

- “Features Nearing End of Life” on page 16
- “Solaris 10 11/06 Operating System” on page 16
- “Solaris Volume Manager GUI” on page 17
- “Loopback File System (LOFS)” on page 17
- “HA-Apache Configuration Wizard” on page 17
- “Accessibility Features for People With Disabilities” on page 17

This section contains information about Sun Cluster compatibility issues, such as features nearing end of life.

- Additional Sun Cluster framework compatibility issues are documented in Chapter 1, “Planning the Sun Cluster Configuration,” in *Sun Cluster Software Installation Guide for Solaris OS*.

- Additional Sun Cluster upgrade compatibility issues are documented in “Upgrade Requirements and Software Support Guidelines” in *Sun Cluster Software Installation Guide for Solaris OS*.
- For other known problems or restrictions, see “[Known Issues and Bugs](#)” on page 25.

Features Nearing End of Life

The following features are nearing end of life in Sun Cluster 3.2 software.

Sun Cluster 3.0

As of the Sun Cluster 3.2 release, Sun Cluster 3.0 is being end of lived. The Sun Cluster 3.0 part number will no longer available.

Solaris 8

As of Sun Cluster 3.2, Sun Cluster will not longer support Solaris 8.

Rolling Upgrade

The rolling upgrade functionality might not be available for upgrading Sun Cluster to the next minor release. In that case, other procedures will be provided that are designed to limit cluster outage during those software upgrades.

Solaris 10 11/06 Operating System

The following known issues might affect the operation of the Sun Cluster 3.2 release with Solaris 10 11/06 operating system. Contact your Sun representative to obtain the necessary patches to fix these issues. For more information, refer to Infodoc 87995.

- | | |
|---------|---|
| 6252216 | metaset command fails after the rpcbind server is restarted. |
| 6331216 | disksets: devid information not written to a newly created diskset. |
| 6345158 | svm exited with error 1 in step cmmstep5, nodes panic. |
| 6367777 | fsck: svc:/system/filesystem/usr fails to start from milestone none. |
| 6401357 | Solaris Volume Manager (SVM) does not show metaset after cluster upgrade in x86 |
| 6402556 | commd timeout should be a percentage of metaclust timeout value. |
| 6474029 | metaset -s diskset -t should take ownership of a cluster node after reboot. |
| 6496941 | SVM still removes the diskset if the Sun Cluster nodeid file is missing. |

- 6367777 fsck* svc:/system/filesystem/usr fails to start from milestone.
- 6367948 New fsck_ufs (1M) has nits when dealing with already mounted file.
- 6425930 Node panics with CMM:cluster lost operational quorum in amd64.
- 6361537 create_ramdisk: cannot seek to offset -1.
- 6393691 Add etc/cluster/nodeid entry to filelist.ramdisk.
- 6344611 create_ramdisk needs to react less poorly to missing files or directories.
- 6462748 devfsadm link removal does not provide full interpose support.

Solaris Volume Manager GUI

The Enhanced Storage module of Solaris Management Console (Solaris Volume Manager) is not compatible with Sun Cluster software. Use the command-line interface or Sun Cluster utilities to configure Solaris Volume Manager software.

Loopback File System (LOFS)

Sun Cluster 3.2 software does not support the use of LOFS under certain conditions. If you must enable LOFS on a cluster node, such as when you configure non-global zones, first determine whether the LOFS restrictions apply to your configuration. See the guidelines in “Solaris OS Feature Restrictions” in *Sun Cluster Software Installation Guide for Solaris OS* for more information about the restrictions and workarounds that permit the use of LOFS when restricting conditions exist.

HA-Apache Configuration Wizard

The Sun Cluster data service configuration wizard for Apache does not support Apache 2.

Accessibility Features for People With Disabilities

To obtain accessibility features that have been released since the publishing of this media, consult Section 508 product assessments that are available from Sun upon request to determine which versions are best suited for deploying accessible solutions.

Commands Modified in This Release

This section describes changes to the Sun Cluster command interfaces that might cause user scripts to fail.

Object-Oriented Command Line Interface

Beginning with the Sun Cluster 3.2 release, Sun Cluster software includes an object-oriented command set. Although Sun Cluster software still supports the original command set, Sun Cluster procedural documentation uses only the object-oriented command set. For more information about the object-oriented command set, see the `Intro(1CL)` man page. For a list of object-oriented commands for common Sun Cluster procedures, see the *Sun Cluster Quick Reference*.

`scinstall` Command

The following options to the `scinstall` command have changed in the Sun Cluster 3.2 release:

- The `-d` option has been removed from use with the `-i` option. The `scinstall` command no longer performs installation of Sun Cluster software packages. Instead, use the `installer` command. See “How to Install Sun Cluster Framework and Data-Service Software Packages” in *Sun Cluster Software Installation Guide for Solaris OS* for more information. The `-d` option is still valid with the `-a`, `-c`, and `-u` options.
- The `-k` option is no longer necessary. It is still provided only for backwards compatibility with user scripts that use this option.
- The `-M` option has been removed from use. Instead, use the appropriate patch management tool for the version of the Solaris OS that your cluster runs. See “[Patches and Required Firmware Levels](#)” on page 38 for more information.

`scconf` Command

The `-q` option of the `scconf` command has been modified to distinguish between shared local quorum devices (SCSI) and other types of quorum devices (including NetApp NAS devices). Use the `name` suboption to specify the name of the attached shared-storage device when adding or removing a shared quorum device to or from the cluster. This suboption can also be used with the `change` form of the command to change the state of a quorum device. The `globaldev` suboption can still be used for SCSI shared-storage devices, but the `name` suboption must be used for all other types of shared storage devices. For more information about this change to

scconf and working with quorum devices, see `scconf(1M)`, `scconf_quorum_dev_netapp_nas(1M)`, `scconf_quorum_dev_netapp_nas(1M)`, and `scconf_quorum_dev_scsi(1M)`.

Resource Properties

It is no longer necessary to modify the `Network_resources_used` resource property directly. Instead, use the `Resource_dependencies` property. The RGM automatically updates the `Network_resources_used` property based on the settings of the `Resource_dependencies` property. For more information about the current uses of these two resource properties, see the `r_properties(5)` man page.

Product Name Changes

This section provides information about product name changes for applications that Sun Cluster software supports. Depending on the Sun Cluster software release that you are running, your Sun Cluster documentation might not reflect the following product name changes.

Current Product Name	Former Product Name
Sun Cluster Manager	SunPlex Manager
Sun Cluster Agent Builder	SunPlex Agent Builder
Sun Java System Application Server	Sun ONE Application Server
Sun Java System Application Server EE (HADB)	Sun Java System HADB
Sun Java System Message Queue	Sun ONE Message Queue
Sun Java System Web Server	<ul style="list-style-type: none"> ■ Sun ONE Web Server ■ iPlanet Web Server ■ Netscape™ HTTP

Supported Products

This section describes the supported software and memory requirements for Sun Cluster 3.2 software.

- **Memory Requirements** – Sun Cluster 3.2 software requires the following memory requirements for every cluster node:
 - Minimum of 512 MB of physical RAM (2GB typical).
 - Minimum of 6GB of available hard drive space.

Actual physical memory and hard drive requirements are determined by the applications installed. Consult the application's documentation or contact the application vendor to calculate additional memory and hard drive requirements.

- **RSMAPI** – Sun Cluster 3.2 software supports the Remote Shared Memory Application Programming Interface (RSMAPI) on RSM-capable interconnects, such as PCI-SCI.
- **Solaris Operating System (OS)** – Sun Cluster 3.2 software and Quorum Server software requires the following minimum versions of the Solaris OS:

- **Solaris 9** – Solaris 9 9/05 SPARC only
- **Solaris 10** – Solaris 10 11/06

- **Solaris Trusted Extensions**

Sun Cluster 3.2 supports Solaris non-global zones within a cluster. Solaris 10 11/06 includes support for Solaris Trusted Extensions. Solaris Trusted Extensions uses non-global zones as well. The interaction between Sun Cluster and Solaris Trusted Extensions using non-global zones has not been tested. Customers are advised to proceed with caution when using these technologies.

- **Volume managers**

Platform	Operating System	Volume Manager	Cluster Feature
SPARC	Solaris 9	Solaris Volume Manager.	Solaris Volume Manager for Sun Cluster.
		VERITAS Volume Manager 4.1. This support requires VxVM 4.1 MP1 patch 124354-02.	VERITAS Volume Manager 4.1 cluster feature.
		VERITAS Volume Manager components that are delivered as part of VERITAS Storage Foundation 4.1. This support requires VxVM 4.1 MP1 patch 124354-02.	VERITAS Volume Manager 4.1 cluster feature.
	VERITAS Volume Manager components that are delivered as part of VERITAS Storage Foundation 5.0. This support requires a VxVM patch. Contact VERITAS for patch number.	N/A - Sun Cluster 3.2 does not support the VxVM cluster feature.	
	Solaris 10	Solaris Volume Manager.	Solaris Volume Manager for Sun Cluster.
		VERITAS Volume Manager 4.1. This support requires VxVM 4.1 MP1 patch 124354-02.	VERITAS Volume Manager 4.1 with cluster feature.
		VERITAS Volume Manager 4.1. This support requires VxVM 4.1 MP1 patch 124354-02.	VERITAS Volume Manager 4.1 with cluster feature.
		VERITAS Volume Manager components that are delivered as part of VERITAS Storage Foundation 5.0. This support requires a VxVM patch. Contact VERITAS for patch number.	N/A - Sun Cluster 3.2 does not support the VxVM cluster feature.
x86	Solaris 10	Solaris Volume Manager.	Solaris Volume Manager for Sun Cluster.
		VERITAS Volume Manager components that are delivered as part of VERITAS Storage Foundation 4.1.	N/A - Sun Cluster 3.2 does not support the VxVM cluster feature on the x86 platform.

- **File systems**
 - **SPARC based systems**

- **On Solaris 9** – Solaris UFS, Sun StorEdge QFS, VERITAS File System 4.1, and VERITAS File System components that are delivered as part of VERITAS Storage Foundation 4.1 and 5.0.
- **On Solaris 10** – Solaris UFS, Solaris ZFS, Sun StorEdge QFS, VERITAS File System 4.1, and VERITAS File System components that are delivered as part of VERITAS Storage Foundation 4.1 and 5.0.
- **x86 based systems**
 - **On Solaris 10** – Solaris UFS.
- **Sun StorEdge™ Availability Suite 10**
- **Sun Management Center 3.6.1**
- **Data services (agents)** – Contact your Sun sales representative for the complete list of supported data services and application versions.

The following Sun Cluster data services support non-global zones:

- Sun Cluster Data Service for Apache
- Sun Cluster Data Service for Apache Tomcat
- Sun Cluster Data Service for DHCP
- Sun Cluster Data Service for Domain Name Service (DNS)
- Sun Cluster Data Service for Kerberos
- Sun Cluster Data Service for MySQL
- Sun Cluster Data Service for N1 Grid Service Provisioning Server
- Sun Cluster Data Service for Oracle
- Sun Cluster Data Service for Oracle Application Server
- Sun Cluster HA for PostgreSQL
- Sun Cluster Data Service for Samba
- Sun Cluster Data Service for Sun Java System Application Server
- Sun Cluster Data Service for Sun Java System Message Queue Server
- Sun Cluster Data Service for Sun Java System Web Server

Note – Procedures for the version of Sun Cluster HA for Sun Java™ System Directory Server that uses Sun Java System Directory Server 5.0 and 5.1 are located in the *Sun Cluster 3.1 Data Service for Sun ONE Directory Server*. For later versions of Sun Java System Directory Server, see the Sun Java System Directory Server product documentation.

The following data services are not supported on Solaris 10 in this Sun Cluster release.

- Sun Cluster Data Service for Agfa IMPAX
- Sun Cluster Data Service for SWIFT Alliance Access
- Sun Cluster Data Service for SWIFT Alliance Gateway

The following is a list of Sun Cluster data services and their resource types.

Data Service	Sun Cluster Resource Type
Sun Cluster HA for Agfa IMPAX	SUNW.gds
Sun Cluster HA for Apache	SUNW.apache
Sun Cluster HA for Apache Tomcat	SUNW.sctomcat
Sun Cluster HA for BroadVision One-To-One Enterprise	SUNW.bv
Sun Cluster HA for DHCP	SUNW.gds
Sun Cluster HA for DNS	SUNW.dns
Sun Cluster HA for MySQL	SUNW.gds
Sun Cluster HA for NetBackup	SUNW.netbackup_master
Sun Cluster HA for NFS	SUNW.nfs
Sun Cluster Oracle Application Server	SUNW.gds
Sun Cluster HA for Oracle E-Business Suite	SUNW.gds
Sun Cluster HA for Oracle	SUNW.oracle_server
	SUNW.oracle_listener
Sun Cluster Support for Oracle Real Application Clusters	SUNW.rac_framework
	SUNW.rac_udlm
	SUNW.rac_svm
	SUNW.rac_cvm
	SUNW.rac_hwraid
	SUNW.oracle_rac_server
	SUNW.oracle_listener
Sun Cluster HA for PostgreSQL	SUNW.gds
Sun Cluster HA for Samba	SUNW.gds
Sun Cluster HA for SAP	SUNW.sap_ci
	SUNW.sap_ci_v2
	SUNW.sap_as
	SUNW.sap_as_v2
Sun Cluster HA for SAP liveCache	SUNW.sap_livecache
	SUNW.sap_xserver

Data Service	Sun Cluster Resource Type
Sun Cluster HA for SAP DB	SUNW.sapdb
	SUNW.sap_xserver
Sun Cluster HA for SAP Web Application Server	SUNW.sapenq
	SUNW.saprepl
	SUNW.sapscs
	SUNW.sapwebas
Sun Cluster HA for Siebel	SUNW.sblgtwy
	SUNW.sblsrvr
Sun Cluster HA for Solaris Containers	SUNW.gds
Sun Cluster HA for N1 Grid Engine	SUNW.gds
Sun Cluster HA for Sun Java System Application Server supported versions before 8.1	SUNW.s1as
Sun Cluster HA for Sun Java System Application Server supported versions as of 8.1	SUNW.jsas
	SUNW.jsas-na
Sun Cluster HA for Sun Java System Application Server EE (supporting HADB versions before 4.4)	SUNW.hadb
Sun Cluster HA for Sun Java System Application Server EE (supporting HADB versions as of 4.4)	SUNW.hadb_ma
Sun Cluster HA for Sun Java System Message Queue	SUNW.s1mq
Sun Cluster HA for Sun Java System Web Server	SUNW.iws
Sun Cluster HA for SWIFTAlliance Access	SUNW.gds
Sun Cluster HA for SWIFTAlliance Gateway	SUNW.gds
Sun Cluster HA for Sybase ASE	SUNW.sybase
Sun Cluster HA for WebLogic Server	SUNW.wls
Sun Cluster HA for WebSphere MQ	SUNW.gds
Sun Cluster HA for WebSphere MQ Integrator	SUNW.gds

Sun Cluster Security Hardening

Sun Cluster Security Hardening uses the Solaris operating system hardening techniques recommended by the Sun BluePrints™ program to achieve basic security hardening for clusters. The Solaris Security Toolkit automates the implementation of Sun Cluster Security Hardening.

The Sun Cluster Security Hardening documentation is available at <http://www.sun.com/blueprints/0203/817-1079.pdf>. You can also access the article from <http://www.sun.com/software/security/blueprints>. From this URL, scroll down to the Architecture heading to locate the article “Securing the Sun Cluster 3.x Software.” The documentation describes how to secure Sun Cluster 3.x deployments in a Solaris environment. The description includes the use of the Solaris Security Toolkit and other best-practice security techniques recommended by Sun security experts. The following data services are supported by Sun Cluster security hardening:

- Sun Cluster HA for Apache
- Sun Cluster HA for Apache Tomcat
- Sun Cluster HA for BEA WebLogic Server
- Sun Cluster HA for DHCP
- Sun Cluster HA for DNS
- Sun Cluster HA for MySQL
- Sun Cluster HA for NetBackup
- Sun Cluster HA for NFS
- Sun Cluster HA for Oracle E-Business Suite
- Sun Cluster HA for Oracle
- Sun Cluster Support for Oracle Real Application Clusters
- Sun Cluster HA for PostgreSQL
- Sun Cluster HA for Samba
- Sun Cluster HA for Siebel
- Sun Cluster HA for Solaris Containers
- Sun Cluster HA for SWIFTAlliance Access
- Sun Cluster HA for Sun Java System Directory Server
- Sun Cluster HA for Sun Java System Message Queue
- Sun Cluster HA for Sun Java System Messaging Server
- Sun Cluster HA for Sun Java System Web Server
- Sun Cluster HA for Sybase ASE
- Sun Cluster HA for WebSphere MQ
- Sun Cluster HA for WebSphere MQ Integrator

Known Issues and Bugs

The following known issues and bugs affect the operation of the Sun Cluster 3.2 release. Bugs are grouped into the following categories:

- “Administration” on page 26
- “Data Services” on page 29
- “Installation” on page 31
- “Localization” on page 33
- “Quorum Server” on page 35
- “Runtime” on page 35

- “Upgrade” on page 37

Administration

New CLI Commands Can Fail to Display Some Localized Characters, Especially European Characters (6466002)

Problem Summary: The new object-oriented CLI commands can fail to display extended characters when executed in some non-US locales. The original CLI is not affected by this issue.

Workaround: Switch to the C (English) locale.

The `clnode remove -f` Option Fails to Remove the Node with the Solaris Volume Manager Device Group (6471834)

Problem Summary: The `clnode remove --force` command should remove nodes from the metaset. The *Sun Cluster System Administration Guide for Solaris OS* provides procedures for removing a node from the cluster. These procedures instruct the user to run the `metaset` command for the Solaris Volume Manager disk set removal prior to running `clnode remove`.

Workaround: If the procedures were not followed, it might be necessary to clear the stale node data from the CCR in the usual way: From an active cluster node, use the `metaset` command to clear the node from the Solaris Volume Manager disk sets. Then run `clnode clear --force obsolete_nodename`.

`scsnapshot` is Nonfunctional With Solaris 10 SUNWCuser Meta Cluster (6477905)

Problem Summary: On a cluster installed with the Solaris 10 End User software group, SUNWCuser, running the `scsnapshot` command might fail with the following error:

```
# scsnapshot -o
...
/usr/cluster/bin/scsnapshot[228]: /usr/perl5/5.6.1/bin/perl: not found
```

Workaround: Do either of the following:

- Install the Solaris Entire Distribution software group.
- Install the following Perl packages: SUNWp15u, SUNWp15v, SUNWp15p.

The `cldevice list -v` Command Can Generate a Core Dump in Some Situations (6485587)

Problem Summary: The `cldevice list -v` command can generate a core dump when the following criteria are both met:

- The verbose (`-v`) option is used.
- There are more than 100 rows to be displayed.

Workaround: Use the `show` subcommand instead of the `list` subcommand. Alternatively, use the original CLI command `scdidadm -L` instead of `cldevice list -v`.

The `clresourcegroup evacuate` Command Can Result in a Delay in Failover Back to the Evacuated Node (6485870)

Problem Summary: If the `-T` (timeout) option is not specified with the `clresourcegroup evacuate` command, the established default is not used. This can cause a delay in failover back to the evacuated node once the node reboots. The `scswitch` and `clnode` commands do not have this problem.

Workaround: When you evacuate a node by using the `clresourcegroup evacuate` command, always specify the `-T` option on the command line. The option `-T 60` will produce the correct default behavior.

Alternately, use the original CLI command `scswitch -S -h node` instead of the `clresourcegroup evacuate -n node` command. Additionally, when this delay in failover back to the evacuated node is observed, use the appropriate `clresourcegroup online`, `clresourcegroup switch`, or `scswitch` command to force a switch of the resource group back to the evacuated node.

Entries in the `AuxnodeList` Property Causes SEGV During Scalable Resource Creation (6494243)

Problem Summary: The `AuxnodeList` property of the shared-address resource cannot be used during shared-address resource creation. This will cause validation errors and SEGV when the scalable resource that depends on this shared address network resource is created. The scalable resource's validate error message is in the following format:

```
Method methodname (scalable svc) on resource resourcename stopped or terminated
due to receipt of signal 11
```

Also, the core file is generated from `ssm_wrapper`. Users will not be able to set the `AuxnodeList` property and thus cannot identify the cluster nodes that can host the shared address but never serve as primary.

Workaround: On one node, re-create the shared-address resource without specifying the `AuxnodeList` property. Then rerun the scalable-resource creation command and use the shared-address resource that you re-created as the network resource.

`clquorumserver` Start and Stop Commands Should Set the Startup State Properly for Next Boot (6496008)

Problem Summary: The Quorum Server command `clquorumserver` does not set the state for the startup mechanism correctly for the next reboot.

Workaround: Perform the following tasks to start or stop Quorum Server software.

How To Start Quorum Server Software on the Solaris 10 OS

1. Display the status of the `quorumserver` service.

```
# svcs -a | grep quorumserver
```

If the service is disabled, output appears similar to the following:

```
disabled          3:33:45  svc:/system/cluster/quorumserver:default
```

2. Start Quorum Server software.

- If the `quorumserver` service is disabled, use the `svcadm enable` command.

```
# svcadm enable svc:/system/cluster/quorumserver:default
```

- If the `quorumserver` service is online, use the `clquorumserver` command.

```
# clquorumserver start +
```

How to Stop Quorum Server Software on the Solaris 10 OS

- Disable the `quorumserver` service.

```
# svcadm disable svc:/system/cluster/quorumserver:default
```

How To Start Quorum Server Software on the Solaris 9 OS

1. Start Quorum Server software.

```
# clquorumserver start +
```

2. Rename the `/etc/rc2.d/.S99quorumserver` file as `/etc/rc2.d/S99quorumserver`.

```
# mv /etc/rc2.d/.S99quorumserver /etc/rc2.d/S99quorumserver
```

How To Stop Quorum Server Software on the Solaris 9 OS

1. Stop Quorum Server software.

```
# clquorumserver stop +
```

2. Rename the `/etc/rc2.d/S99quorumserver` file as `/etc/rc2.d/.S99quorumserver`.

```
# mv /etc/rc2.d/S99quorumserver /etc/rc2.d/.S99quorumserver
```

Data Services

Creation of Node Agent Resource for Sun Cluster HA for Sun Java Systems Application Server Succeeds Even if Resource Dependency is Not Set on Domain Administration Server (DAS) Resource (6262459)

Problem Summary: When creating the node agent (NA) resource in Sun Cluster HA for Application Server, the resource gets created even if there is no dependency set on the DAS resource. The command should error out if the dependency is not set, because a DAS resource must be online in order to start the NA resource.

Workaround: While creating the NA resource, make sure you set a resource dependency on the DAS resource.

Oracle Resources Fail to Stay Offline When a Depended on Device Group Becomes Disabled (6479501)

Problem Summary: When a fault occurs in a disk group in an Oracle RAC database configuration, the `SUNW.ScalDeviceGroup` resource that manages the disk group becomes disabled, but the Oracle resources go offline and then online again. The Oracle resources and the database instances encounter errors while trying to come online, due to the unavailability of the disk group. These errors might include a core dump of the resource processes or of the Oracle instances, and error messages in the associated log files. A similar outcome results if a `SUNW.ScalMountPoint` resource that manages a file system mount point becomes disabled.

The underlying cause of this problem is that the offline-restart resource dependency does not work properly if the dependent resource has a `Postnet_stop` method. When the depended-on resource goes offline, the dependent resource executes the `Stop` method, immediately followed by the `Start` method. Instead, it should execute `Stop` and `Postnet_stop` to go offline.

Workaround: Manually disable the Oracle resources, and any other dependent resources such as `ScaMountPoint`, until the disk problem is repaired. After the disk group is repaired and verified from the volume manager point of view, re-enable the disk group resource; then re-enable the other resources that you disabled manually. In the case of `ScaMountPoint` being disabled, after the file system issue is resolved and verified by a successful mount command, re-enable the mount point resource; then re-enable the other resources that you disabled manually.

Additional Prerequisite for Using the Oracle 10g RAC Configuration Wizard to Create Proxy Resources Within Oracle CRS (6488017)

Problem Summary: The Oracle 10g RAC configuration wizard fails to create proxy resources in CRS if the corresponding Oracle database instances are not enabled and online. This can be confirmed by the following the error output of the commands executed by the wizard:

CRS-0200: CRS general failure. No reason is provided.

Workaround: Run the following Oracle commands to bring the database instances online before running the create step in the Oracle RAC configuration wizard:

```
# srvctl enable database -d DB_Name
# srvctl start database -d DN_Name
# srvctl status database -d DB_Name
```

The output should look similar to the following:

```
Instance DB_INSTANCE_1 is running on node node1
Instance DB_INSTANCE_2 is running on node node2
```

SAP Central Instance Resource Enters stop_failed State After Being Disconnected From Public Network (6498935)

Problem Summary: IPv6 plumbing on the interconnects, which is required for forwarding of IPv6 scalable service packets, will no longer be enabled by default. The IPv6 interfaces, as seen when using the `ifconfig` command, will no longer be plumbed on the interconnect adapters by default.

Workaround: Manually enable IPv6 scalable service support.

▼ How to Manually Enable IPv6 Scalable Service Support

Before You Begin

Ensure that you have prepared all cluster nodes to run IPv6 services. These tasks include proper configuration of network interfaces, server/client application software, name services, and routing infrastructure. Failure to do so might result in unexpected failures of network applications. For more information, see your Solaris system-administration documentation for IPv6 services.

1 # On each node, add the following entry to the `/etc/system` file.

```
set cl_comm:ifk_disable_v6=0
```

2 On each node, enable IPv6 plumbing on the interconnect adapters.

```
# /usr/cluster/lib/sc/config_ipv6
```

The `config_ipv6` utility brings up an IPv6 interface on all cluster interconnect adapters that have a link-local address. The utility enables proper forwarding of IPv6 scalable service packets over the interconnects.

Alternately, you can reboot each cluster node to activate the configuration change.

Installation

Autodiscovery With InfiniBand Configurations Can Sometimes Suggest Two Paths Using the Same Adapter (6299097)

Problem Summary: If InfiniBand is used as the cluster transport and there are two adapters on each node with two ports per adapter and a total of two switches, the `scinstall` utility's adapter autodiscovery could suggest two transport paths that use the same adapter.

Workaround: Manually specify the transport adapters on each node.

IPv6 Scalable Service Support is Not Enabled by Default (6332656)

Problem Summary: If InfiniBand is used as the cluster transport and there are two adapters on each node with two ports per adapter and a total of two switches, the `scinstall` utility's adapter autodiscovery could suggest two transport paths that use the same adapter.

Workaround: Manually specify the transport adapters on each node.

IPv6 Scalable Service Support is Not Enabled by Default (6332656)

Problem Summary: IPv6 plumbing on the interconnects, which is required for forwarding of IPv6 scalable service packets, will no longer be enabled by default. The IPv6 interfaces, as seen when using the `ifconfig` command, will no longer be plumbed on the interconnect adapters by default.

Workaround: Manually enable IPv6 scalable service support.

▼ How to Manually Enable IPv6 Scalable Service Support

Before You Begin

Ensure that you have prepared all cluster nodes to run IPv6 services. These tasks include proper configuration of network interfaces, server/client application software, name services, and routing infrastructure. Failure to do so might result in unexpected failures of network applications. For more information, see your Solaris system-administration documentation for IPv6 services.

- 1 On each node, add the following entry to the `/etc/system` file.

```
set cl_comm:ifk_disable_v6=0
```

- 2 On each node, enable IPv6 plumbing on the interconnect adapters.

```
# /usr/cluster/lib/sc/config_ipv6
```

The `config_ipv6` utility brings up an IPv6 interface on all cluster interconnect adapters that have a link-local address. The utility enables proper forwarding of IPv6 scalable service packets over the interconnects.

Alternately, you can reboot each cluster node to activate the configuration change.

`clnode add` **Fails to Add a Node from an XML File if the File Contains Direct-Connect Transport Information (6485249)**

Problem Summary: If the `clnode add` command is attempted using an XML file on a that is using direct-connect transport, the command misinterprets the cable information and adds the wrong configuration information. As a result, the joining node is not able to join the cluster.

Workaround: Use the `scinstall` command to add a node to the cluster when the cluster transport is direct connected.

`cluster create` **Command Does Not Properly Determine a Remote Access Method (6482672)**

Problem Summary: The `cluster create` command has a dependency on being able to remotely execute commands on other cluster nodes. The command should automatically attempt to test for each of the following three methods:

- A special Sun Cluster remote access method which can be selected in the `installer` utility at the time that each node is installed.
- Access using `ssh`.
- Access using `rsh`.

The first method is the least secure and it is not selected at install time, so it is not used. Furthermore, the first method is automatically disabled once the cluster is configured for the first time. And, if `cluster create` is used to recreate a cluster which was created and later dismantled, this first method will still be disabled. The defect is that the `cluster create` command only recognizes the first method.

Workaround: Choose one of the following options, depending on your security requirements:

- If the more secure `ssh` or `rsh` option is required, use the `scinstall` utility or command to create or re-create a new cluster.
- If the built-in remote-configuration option is acceptable, enable the `cluster create` command by creating the file `/etc/cluster/remotefconfiguration` on each of the new cluster nodes:

```
# touch /etc/cluster/remotefconfiguration
```

Note – If the `cluster create` command is being used for the first time on a set of nodes that is installed with the `installer` utility, and the installer's remote-configuration option was not deselected, you do not need to perform this workaround.

The /etc/nsswitch.conf File is Not Updated with host and netmasks Database Information During Non-Global Zone Installation (6345227)

Problem Summary: The `scinstall` command updates the `/etc/nsswitch.conf` file to add the `clusterentry` for the `hosts` and `netmasks` databases. This change updates the `/net/nsswitch.conf` file for the global zone. But when a non-global zone is created and installed, the non-global zone receives its own copy of the `/etc/nsswitch.conf` file. The `/etc/nsswitch.conf` files on the non-global zones will not have the `clusterentry` for the `hosts` and `netmasks` databases. Any attempt to resolve cluster-specific private hostnames and IP addresses from within a non-global zone by using `getxbyy` queries will fail.

Workaround: Manually update the `/etc/nsswitchconf` file for non-global zones with the `clusterentry` for the `hosts` and `netmasks` database. This ensures that the cluster-specific private-hostname and IP-address resolutions are available within non-global zones.

Localization

When a Browser is Launched on the French Locale, Sun Cluster Manager Does Not Work Correctly (6450880)

Problem Summary: When the user accesses Sun Cluster Manager from a browser that is launched on the French locale, the user cannot administer resources and resource groups. Sun Cluster Manager does not activate action-item lists, even when resources and resource groups are correctly selected.

Workaround: Access Sun Cluster Manager from browser that is launched on the C (English) locale.

Invalid Number is Reported for the Number of Cluster Node and Private Networks at Changes of Private Network Capability (6481119)

Problem Summary: When using the `scinstall` utility to change the private network configuration, the screens for configuring the netmask report invalid minus or very large numbers for the number of cluster nodes and private networks.

Workaround: Run the `scinstall` utility on the C locale by setting `LANG` or `LC_ALL` to C.

Translated Messages for Quorum Server are Delivered in a Separate Package From Core Translation Packages (6482813)

Problem Summary: Translated messages for the Quorum Server administration programs, such as `clquorumserver`, are delivered in a separate package from the core translation packages. As a result, Quorum Server messages appear only in English if the Quorum Server

packages are not installed on the host. Also, the Japanese language man page is not available on the host. This problem is not observed if all other Sun Cluster packages are installed on the host where Quorum Server software is installed.

Workaround: Install the following packages on the host where Quorum Server software is installed:

- SUNWcsc (Simplified Chinese)
- SUNWdsc (German)
- SUNWesc (Spanish)
- SUNWfsc (French)
- SUNWhsc (Traditional Chinese)
- SUNWjsc (Japanese)
- SUNWksc (Korean)

If the Japanese man page is needed on the quorum server, install the SUNWj_scm_{an} (Japanese man page) package.

Localized Messages in `clsetup`, `scsetup`, and `scinstall` Show Badly Aligned Messages on Some Localized Environments (6492156)

Problem Summary: On some localized environments like UTF-8, the `clsetup`, `scsetup`, and `scinstall` commands show misaligned messages when messages have long single lines.

Workaround: Perform the following steps:

1. Save a backup copy of the `usr/cluster/lib/scadmin/lib/sc_common` file as `usr/cluster/lib/scadmin/lib/sc_common.org`.
2. In the `usr/cluster/lib/scadmin/lib/sc_common` file, remove the `b` option from the `fold` command in the `sc_print_para()` function.

Original line:

```
echo "${text}" | fold -bs -w 70 | sed 's/^/ /'
```

Revised line:

```
echo "${text}" | fold -s -w 70 | sed 's/^/ /'
```

Setting up Telemetry Fails in Localized Environment (6493402)

Problem Summary: If you are setting up telemetry attributes by using the `scsetup` or `clsetup` utility in a locale other than the C (English) locale, you might encounter a failure. The translated error message is similar to Error: cannot retrieve the name of the cluster node hosting the database on which system resource monitoring relies.

Workaround: Use the C locale when you set up telemetry attributes by using the `scsetup` or `clsetup` utility.

Installer Displays Incorrect Swap Size for the Sun Cluster 3.2 Simplified Chinese Version (6495984)

Problem Summary: The Sun Cluster 3.2 installer displays a warning message about short swap when installing the Sun Cluster 3.2 Simplified Chinese version of the software. The installer provides an incorrect swap size of 0.0KB size on the system requirements check screen.

Workaround: If the swap size is larger than the system requirement, you can safely ignore this problem. The SC 3.2 installer on the C or English locale can be used for installation and this version checks swap size correctly.

Quorum Server

Cluster Nodes Fail to Contact Quorum Server at Boot (6449588)

Problem Summary: On some x86 clusters, the administrator might successfully configure a quorum server that is on a different network than the cluster. At boot time, the cluster nodes fail to contact the quorum server.

Workaround: Configure a quorum server that is on the same network as the cluster nodes.

Runtime

SAP cleanipc Binary Needs User_env Parameter for LD_Library_Path

Problem Summary: The cleanipc fails if the runtime linking environment does not contain the /sapmnt/SAPSID/exe path.

Workaround: As the Solaris root user, add the /sapmnt/SAPSID/exe path to the default library in ld.config file.

To configure the runtime linking environment default library path for 32-bit applications, enter the following command:

```
# crle -u -l 32 bit SAP library path
```

To configure the runtime linking environment default library path for 64-bit applications, enter the following command:

```
# crle -64 -u -l 64 bit SAP library path
```

Node Panics Due to a `metac lust` Return Step Error: RPC: Program Not Registered (6256220)

Problem Summary: When a cluster shutdown is performed, the UCMMD can go into a reconfiguration on one or more of the nodes if one of the nodes leaves the cluster slightly ahead of the UCMMD. When this occurs, the shutdown stops the `rpc .md` command on the node while the UCMMD is trying to perform the return step. In the return step, the `metac lust` command gets an RPC timeout and exits the step with an error, due to the missing `rpc .mdcmd` process. This error causes the UCMMD to abort the node, which might cause the node to panic.

Workaround: You can safely ignore this problem. When the node boots back up, Sun Cluster software detects this condition and allows the UCMMD to start, despite the fact that an error occurred in the previous reconfiguration.

Sun Cluster Resource Validation Does Not Accept the Hostname for IPMP Groups for the `netiflist` Property (6383994)

Problem Summary: Sun Cluster resource validation does not accept the hostname for IPMP groups for the `netiflist` property during logical-hostname or shared-address resource creation.

Workaround: Use the node name instead of the node ID when you specify the IPMP group names during logical-hostname and shared-address resource creation.

Node Does Not Reboot if it has Local Disks under Disk-Path Monitoring, Even if All Paths to Shared Storage Fail (6479327)

Problem Summary: By default, the `scdpmd` command monitors all disks, including local disks, in the cluster. The `scdpmd` command reboots a node only if the following criteria are all met:

- All disks, including local disks, have failed on that node.
- Another node has access to at least one of the failed disks.
- The `reboot_on_path_failure` flag is enabled on that node.

But the `scdpmd` command will not reboot a node if the path to a local disk on the node is still accessible, even if the paths to all shared disks on the node have failed and even if the `reboot_on_path_failure` is enabled for that node.

Workaround: To avoid this problem, unmonitor all local disk paths in the cluster by typing the following, where *nodename* is the name of the node where you want to unmonitor local disk paths:

```
# /usr/cluster/bin/scdpm -u nodename:localdiskpath
```

The vxlufinish Script Returns an Error When the Root Disk is Encapsulated (6448341)

Problem Summary: This problem is seen when the original disk is root encapsulated and a live upgrade is attempted from VxVM 3.5 on Solaris 9 8/03 OS to VxVM 5.0 on Solaris 10 6/06 OS. The vxlufinish script fails with the following error.

```
#./vslufinish -u 5.10

VERITAS Volume Manager VxVM 5.0
Live Upgrade finish on the Solairs release <5.10>

Enter the name of the alternate root diskgroup: altrootdg
ld.so.1: vxparms: fatal: libvxscsi.so: open failed: No such file or directory
ld.so.1: vxparms: fatal: libvxscsi.so: open failed: No such file or directory
Killed
ld.so.1: ugettext: fatal: libvxscsi.so: open failed: No such file or directory
ERROR:vxlufinish Failed: /altroot.5.10/usr/lib/vxvm/bin/vxencap -d -C 10176
-c -p 5555 -g
-g altrootdg rootdisk=c0t1d0s2
Please install, if 5.0 or higher version of VxVM is not installed
on alternate bootdisk.
```

Workaround: Use the standard upgrade or dual-partition upgrade method instead.

Contact Sun support or your Sun representative to learn whether Sun Cluster 3.2 Live Upgrade support for VxVM 5.0 becomes available at a later date.

Upgrade

The vxlustart Script Fails to Create the Alternate Boot Environment During a Live Upgrade (6445430)

Problem Summary: This problem is seen when upgrading VERITAS Volume Manager (VxVM) during a Sun Cluster live upgrade. The vxlustart script is used to upgrade the Solaris OS and VxVM from the previous version. The script fails with error messages similar to the following:

```
# ./vxlustart -u 5.10 -d c0t1d0 -s OSimage

VERITAS Volume Manager VxVM 5.0.
Live Upgrade is now upgrading from 5.9 to <5.10>
...
ERROR: Unable to copy file systems from boot environment &lt;source.8876> to BE &lt;dest.8876>.
ERROR: Unable to populate file systems on boot environment &lt;dest.8876>.
```

```
ERROR: Cannot make file systems for boot environment &lt;dest.8876>.  
ERROR: vxlustart: Failed: lucreate -c source.8876 -C /dev/dsk/c0t0d0s2  
-m -:/dev/dsk/c0t1d0s1:swap -m /:/dev/dsk/c0t1d0s0:ufs  
-m /globaldevices:/dev/dsk/c0t1d0s3:ufs -m /mc_metadb:/dev/dsk/c0t1d0s7:ufs  
-m /space:/dev/dsk/c0t1d0s4:ufs -n dest.8876
```

Workaround: Use the standard upgrade or dual-partition upgrade method if you are upgrading the cluster to VxVM 5.0.

Contact Sun support or your Sun representative to learn whether Sun Cluster 3.2 Live Upgrade support for VxVM 5.0 becomes available at a later date.

vxio Major Numbers Different Across the Nodes When the Root Disk is Encapsulated (6445917)

Problem Summary: For clusters that run VERITAS Volume Manager (VxVM), a standard upgrade or dual-partition upgrade of any of the following software fails if the root disk is encapsulated:

- Upgrading the Solaris OS to a different version
- Upgrading VxVM
- Upgrading Sun Cluster software

The cluster node panics and fails to boot after upgrade. This is due to the major-number or minor-number changes made by VxVM during the upgrade.

Workaround: Unencapsulate the root disk before you begin the upgrade.



Caution – If the above procedure is not followed correctly, you may experience serious unexpected problems on all nodes being upgraded. Also, unencapsulation and encapsulation of root disk causes an additional reboot (each time) of the node automatically, increasing the number of required reboots during upgrade.

Patches and Required Firmware Levels

This section provides information about patches for Sun Cluster configurations. If you are upgrading to Sun Cluster 3.2 software, see Chapter 8, “Upgrading Sun Cluster Software,” in *Sun Cluster Software Installation Guide for Solaris OS*. Applying a Sun Cluster 3.2 Core patch does not provide the same result as upgrading the software to the Sun Cluster 3.2 release.

Note – You must be a registered SunSolve™ user to view and download the required patches for the Sun Cluster product. If you do not have a SunSolve account, contact your Sun service representative or sales engineer, or register online at <http://sunsolve.sun.com>.

Patch Management Tools

The PatchPro patch management technology is now available as Patch Manager 2.0 for Solaris 9 OS and as Sun Update Connection 1.0 for Solaris 10 OS.

- **Solaris 9** - Sun Patch Manager 2.0 is available for free download from SunSolve at <http://www.sun.com/software/download/products/40c8c2ad.html>. Documentation for Sun Patch Manager is available at <http://docs.sun.com/app/docs/coll/1152.1>.
- **Solaris 10** - Sun Update Connection is available as patch ID 121118-05 (SPARC) or 121119-05 (x86) or as a download from SunSolve. See <http://www.sun.com/service/sunupdate/gettingstarted.html> for details. Documentation for Sun Update Connection is available at <http://docs.sun.com/app/docs/coll/1320.2>.

Additional information about all patch management options for the Solaris 10 OS is available at <http://www.sun.com/service/sunupdate/>. Additional information for using the Sun patch management tools is provided in the *Solaris Administration Guide: Basic Administration* at <http://docs.sun.com> (<http://docs.sun.com>). Refer to the version of this manual that is published for the Solaris OS release that you have installed.

If some patches must be applied when the node is in noncluster mode, you can apply them in a rolling fashion, one node at a time, unless a patch's instructions require that you shut down the entire cluster. Follow procedures in “How to Apply a Rebooting Patch (Node)” in *Sun Cluster System Administration Guide for Solaris OS* to prepare the node and boot it into noncluster mode. For ease of installation, consider applying all patches at once to a node that you place in noncluster mode.

SunSolve Online

The SunSolve Online Web site provides 24-hour access to the most up-to-date information regarding patches, software, and firmware for Sun products. Access the SunSolve Online site at <http://sunsolve.sun.com> for the most current matrixes of supported software, firmware, and patch revisions.

Sun Cluster 3.2 third-party patch information is provided through a SunSolve Info Doc. This Info Doc page provides any third-party patch information for specific hardware that you intend to use in a Sun Cluster 3.2 environment. To locate this Info Doc, log on to SunSolve. From the SunSolve home page, type **Sun Cluster 3.x Third-Party Patches** in the search criteria box.

Before you install Sun Cluster 3.2 software and apply patches to a cluster component (Solaris OS, Sun Cluster software, volume manager software, data services software, or disk hardware), review each README file that accompanies the patches that you retrieved. All cluster nodes must have the same patch level for proper cluster operation.

For specific patch procedures and tips on administering patches, see Chapter 10, “Patching Sun Cluster Software and Firmware,” in *Sun Cluster System Administration Guide for Solaris OS*.

Sun Cluster 3.2 Documentation

The Sun Cluster 3.2 user documentation set consists of the following collections:

- “Searching Sun Product Documentation” on page 40
- “Sun Cluster 3.2 Software Manuals for Solaris OS” on page 41
- “Sun Cluster 3.2 Data Service Manuals for Solaris OS (SPARC Platform Edition)” on page 41
- “Sun Cluster 3.2 Data Service Manuals for Solaris OS (x86 Platform Edition)” on page 43
- “Sun Cluster 3.1 – 3.2 Hardware Collection for Solaris OS (SPARC Platform Edition)” on page 44
- “Sun Cluster 3.1 — 3.2 Hardware Collection for Solaris OS (x86 Platform Edition)” on page 45

The Sun Cluster 3.2 user documentation is available in PDF and HTML format at the following web site:

<http://docs.sun.com>

Note – Beginning with Sun Cluster 3.2, documentation for individual data services will not be translated. Documentation for individual data services will be available only in English.

Searching Sun Product Documentation

Besides searching for Sun production documentation from the docs.sun.com web site, you can use a search engine of your choice by typing the following syntax in the search field:

```
search-term site:docs.sun.com
```

For example, to search for “broker,” type the following:

```
broker site:docs.sun.com
```

To include other Sun web sites in your search (for example, java.sun.com, www.sun.com, developers.sun.com), use “sun.com” in place of docs.sun.com” in the search field.

Sun Cluster 3.2 Software Manuals for Solaris OS

TABLE 1 Sun Cluster 3.2 Software Collection for Solaris OS Software Manuals

Part Number	Book Title
819-2969	<i>Sun Cluster Concepts Guide for Solaris OS</i>
819-2968	<i>Sun Cluster Overview for Solaris OS</i>
819-2970	<i>Sun Cluster Software Installation Guide for Solaris OS</i>
819-0912	<i>Sun Cluster Quick Start Guide for Solaris OS</i>
819-2971	<i>Sun Cluster System Administration Guide for Solaris OS</i>
819-2972	<i>Sun Cluster Data Services Developer's Guide for Solaris OS</i>
819-2973	<i>Sun Cluster Error Messages Guide for Solaris OS</i>
819-3055	<i>Sun Cluster Reference Manual for Solaris OS</i>
819-2974	<i>Sun Cluster Data Services Planning and Administration Guide for Solaris OS</i>

Sun Cluster 3.2 Data Service Manuals for Solaris OS (SPARC Platform Edition)

TABLE 2 Sun Cluster 3.2 Software Collection for Solaris OS (SPARC Platform Edition): Individual Data Service Manuals

Part Number	Book Title
819-3056	<i>Sun Cluster Data Service for Agfa IMPAX Guide for Solaris OS</i>
819-2975	<i>Sun Cluster Data Service for Apache Guide for Solaris OS</i>
819-3057	<i>Sun Cluster Data Service for Apache Tomcat Guide for Solaris OS</i>
819-0691	<i>Sun Cluster Data Service for BroadVision One-To-One Enterprise Guide for Solaris OS</i>
819-3058	<i>Sun Cluster Data Service for DHCP Guide for Solaris OS</i>
819-2977	<i>Sun Cluster Data Service for DNS Guide for Solaris OS</i>
819-5415	<i>Sun Cluster Data Services for Kerberos Guide for Solaris OS</i>
819-3059	<i>Sun Cluster Data Service for MySQL Guide for Solaris OS</i>
819-3060	<i>Sun Cluster Data Service for N1 Grid Service Provisioning System for Solaris OS</i>
819-0693	<i>Sun Cluster Data Service for NetBackup Guide for Solaris OS</i>

TABLE 2 Sun Cluster 3.2 Software Collection for Solaris OS (SPARC Platform Edition): Individual Data Service Manuals (Continued)

Part Number	Book Title
819-2979	<i>Sun Cluster Data Service for NFS Guide for Solaris OS</i>
819-3061	<i>Sun Cluster Data Service for Oracle Application Server Guide for Solaris OS</i>
819-3062	<i>Sun Cluster Data Service for Oracle E-Business Suite Guide for Solaris OS</i>
819-2980	<i>Sun Cluster Data Service for Oracle Guide for Solaris OS</i>
819-2981	<i>Sun Cluster Data Service for Oracle RAC Guide for Solaris OS</i>
819-5578	<i>Sun Cluster Data Service for PostgreSQL Guide for Solaris OS</i>
819-3063	<i>Sun Cluster Data Service for Samba Guide for Solaris OS</i>
819-2982	<i>Sun Cluster Data Service for MaxDB Guide for Solaris OS</i>
819-2983	<i>Sun Cluster Data Service for SAP Guide for Solaris OS</i>
819-2984	<i>Sun Cluster Data Service for SAP liveCache Guide for Solaris OS</i>
819-2985	<i>Sun Cluster Data Service for SAP Web Application Server Guide for Solaris OS</i>
819-2986	<i>Sun Cluster Data Service for Siebel Guide for Solaris OS</i>
819-3069	<i>Sun Cluster Data Service for Solaris Containers Guide</i>
819-3064	<i>Sun Cluster Data Service for Sun Grid Engine Guide for Solaris OS</i>
819-2988	<i>Sun Cluster Data Service for Sun Java System Application Server Guide for Solaris OS</i>
819-2987	<i>Sun Cluster Data Service for Sun Java System Application Server EE (HADB) Guide for Solaris OS</i>
819-2989	<i>Sun Cluster Data Service for Sun Java System Message Queue Guide for Solaris OS</i>
819-2990	<i>Sun Cluster Data Service for Sun Java System Web Server Guide for Solaris OS</i>
819-3065	<i>Sun Cluster Data Service for SWIFTAlliance Access Guide for Solaris OS</i>
819-3066	<i>Sun Cluster Data Service for SWIFTAlliance Gateway Guide for Solaris OS</i>
819-2991	<i>Sun Cluster Data Service for Sybase ASE Guide for Solaris OS</i>
819-2992	<i>Sun Cluster Data Service for WebLogic Server Guide for Solaris OS</i>
819-3068	<i>Sun Cluster Data Service for WebSphere Message Broker Guide for Solaris OS</i>
819-3067	<i>Sun Cluster Data Service for WebSphere MQ Guide for Solaris OS</i>

Sun Cluster 3.2 Data Service Manuals for Solaris OS (x86 Platform Edition)

TABLE 3 Sun Cluster 3.2 Software Collection for Solaris OS (x86 Platform Edition): Individual Data Service Manuals

Part Number	Book Title
819-2975	<i>Sun Cluster Data Service for Apache Guide for Solaris OS</i>
819-2975	<i>Sun Cluster Data Service for Apache Tomcat Guide for Solaris OS</i>
819-3058	<i>Sun Cluster Data Service for DHCP Guide for Solaris OS</i>
819-2977	<i>Sun Cluster Data Service for DNS Guide for Solaris OS</i>
819-5415	<i>Sun Cluster Data Services for Kerberos Guide for Solaris OS</i>
819-2982	<i>Sun Cluster Data Service for MaxDB Guide for Solaris OS</i>
819-3059	<i>Sun Cluster Data Service for MySQL Guide for Solaris OS</i>
819-3060	<i>Sun Cluster Data Service for N1 Grid Service Provisioning System for Solaris OS</i>
819-2979	<i>Sun Cluster Data Service for NFS Guide for Solaris OS</i>
819-3061	<i>Sun Cluster Data Service for Oracle Application Server Guide for Solaris OS</i>
819-2980	<i>Sun Cluster Data Service for Oracle Guide for Solaris OS</i>
819-2981	<i>Sun Cluster Data Service for Oracle RAC Guide for Solaris OS</i>
819-5578	<i>Sun Cluster Data Service for PostgreSQL Guide for Solaris OS</i>
819-3063	<i>Sun Cluster Data Service for Samba Guide for Solaris OS</i>
819-2985	<i>Sun Cluster Data Service for SAP Web Application Server Guide for Solaris OS</i>
819-3069	<i>Sun Cluster Data Service for Solaris Containers Guide</i>
819-3064	<i>Sun Cluster Data Service for Sun Grid Engine Guide for Solaris OS</i>
819-2987	<i>Sun Cluster Data Service for Sun Java System Application Server EE (HADB) Guide for Solaris OS</i>
819-2988	<i>Sun Cluster Data Service for Sun Java System Application Server Guide for Solaris OS</i>
819-2989	<i>Sun Cluster Data Service for Sun Java System Message Queue Guide for Solaris OS</i>
819-2990	<i>Sun Cluster Data Service for Sun Java System Web Server Guide for Solaris OS</i>
819-2992	<i>Sun Cluster Data Service for WebLogic Server Guide for Solaris OS</i>
819-3067	<i>Sun Cluster Data Service for WebSphere MQ Guide for Solaris OS</i>

TABLE 3 Sun Cluster 3.2 Software Collection for Solaris OS (x86 Platform Edition): Individual Data Service Manuals (Continued)

Part Number	Book Title
819-3068	<i>Sun Cluster Data Service for WebSphere Message Broker Guide for Solaris OS</i>

Sun Cluster 3.1 – 3.2 Hardware Collection for Solaris OS (SPARC Platform Edition)

TABLE 4 Sun Cluster 3.1 - 3.2 Hardware Collection for Solaris OS (SPARC Platform Edition)

Part Number	Book Title
819-2993	<i>Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS</i>
819-2995	<i>Sun Cluster 3.1 - 3.2 With SCSI JBOD Storage Device Manual for Solaris OS</i>
819-3015	<i>Sun Cluster 3.1 - 3.2 With Sun StorEdge 3310 or 3320 SCSI RAID Array Manual for Solaris OS</i>
819-3016	<i>Sun Cluster 3.1 - 3.2 With Sun StorEdge 3510 or 3511 FC RAID Array Manual for Solaris OS</i>
819-3017	<i>Sun Cluster 3.1 - 3.2 With Sun StorEdge 3900 Series or Sun StorEdge 6900 Series System Manual</i>
819-3018	<i>Sun Cluster 3.1 - 3.2 With Sun StorEdge 6120 Array Manual for Solaris OS</i>
819-3020	<i>Sun Cluster 3.1 - 3.2 With Sun StorEdge 6320 System Manual for Solaris OS</i>
819-3021	<i>Sun Cluster 3.1 - 3.2 With Sun StorEdge 9900 Series Storage Device Manual for Solaris OS</i>
819-2996	<i>Sun Cluster 3.1 - 3.2 With StorEdge A1000 Array, Netra st A1000 Array, or StorEdge A3500 System Manual</i>
819-3022	<i>Sun Cluster 3.1 - 3.2 With Sun StorEdge A3500FC System Manual for Solaris OS</i>
819-2994	<i>Sun Cluster 3.1 - 3.2 With Fibre Channel JBOD Storage Device Manual</i>
817-5681	<i>Sun Cluster 3.1 - 3.2 With SCSI JBOD Storage Device Manual for Solaris OS</i>
819-3023	<i>Sun Cluster 3.1 - 3.2 With Sun StorEdge T3 or T3+ Array Manual for Solaris OS</i>
819-3019	<i>Sun Cluster 3.1 - 3.2 With Sun StorEdge 6130 Array Manual</i>
819-3024	<i>Sun Cluster 3.1 - 3.2 With Network-Attached Storage Devices Manual for Solaris OS</i>

Sun Cluster 3.1 — 3.2 Hardware Collection for Solaris OS (x86 Platform Edition)

TABLE 5 Sun Cluster 3.1 – 3.2 Hardware Collection for Solaris OS (x86 Platform Edition)

Part Number	Book Title
819-2993	<i>Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS</i>
817-0180	<i>Sun Cluster 3.1 - 3.2 With Sun StorEdge 3310 or 3320 SCSI RAID Array Manual for Solaris OS</i>
819-3024	<i>Sun Cluster 3.1 - 3.2 With Network-Attached Storage Devices Manual for Solaris OS</i>
819-3021	<i>Sun Cluster 3.1 - 3.2 With Sun StorEdge 9900 Series Storage Device Manual for Solaris OS</i>
819-3020	<i>Sun Cluster 3.1 - 3.2 With Sun StorEdge 6320 System Manual for Solaris OS</i>
819-3019	<i>Sun Cluster 3.1 - 3.2 With Sun StorEdge 6130 Array Manual</i>
819-3018	<i>Sun Cluster 3.1 - 3.2 With Sun StorEdge 6120 Array Manual for Solaris OS</i>
819-3016	<i>Sun Cluster 3.1 - 3.2 With Sun StorEdge 3510 or 3511 FC RAID Array Manual for Solaris OS</i>
819-2995	<i>Sun Cluster 3.1 - 3.2 With SCSI JBOD Storage Device Manual for Solaris OS</i>

Documentation Issues

This section discusses errors or omissions for documentation, online help, or man pages in the Sun Cluster 3.2 release.

- “Software Installation Guide” on page 45
- “Sun Cluster Data Services Planning and Administration Guide” on page 50
- “System Administration Guide” on page 50
- “Data Services Developer's Guide” on page 52
- “Data Service for SAP Guide” on page 53
- “Quorum Server User's Guide” on page 56
- “Man Pages” on page 56

Software Installation Guide

This section discusses errors and omissions in the *Sun Cluster Software Installation Guide for Solaris OS*.

Incorrect Path to the Common Agent Container Security Directory

There is an error in the procedure “How to Finish Upgrade to Sun Cluster 3.2 Software” in *Sun Cluster Software Installation Guide for Solaris OS*. Step 1d gives the path `/etc/opt/SUNWcacao/security`. The correct path is `/etc/cacao/SUNWcacao/security`.

Sun Cluster Module to Sun Management Center No Longer Provides Online Help

The “See Also” section of “How to Load the Sun Cluster Module” in *Sun Cluster Software Installation Guide for Solaris OS* describes how to access online help for the Sun Cluster module to Sun Management Center. This information is no longer valid for the Sun Cluster 3.2 release. The version of Sun Management Center that Sun Cluster 3.2 software supports no longer provides or supports online help. For information about how to navigate within Sun Management Center windows, see the Sun Management Center documentation. For definitions of cluster objects that the Sun Cluster module monitors, see the *Sun Cluster Overview for Solaris OS* or *Sun Cluster Concepts Guide for Solaris OS*.

QFS Incorrectly Included for Use as a Cluster File System

In “Choosing Mount Options for Cluster File Systems” in *Sun Cluster Software Installation Guide for Solaris OS* and “How to Create Cluster File Systems” in *Sun Cluster Software Installation Guide for Solaris OS*, the inclusion of information to configure a cluster file system with QFS is in error. You should only use UFS or VxFS file systems to create a cluster file system. To configure a shared QFS file system to support Oracle RAC, see *Sun Cluster Data Service for Oracle RAC Guide for Solaris OS*.

Incorrect Command to Disable HADB Resource Before Upgrade

In “How to Prepare the Cluster for Upgrade (Dual-Partition)” in *Sun Cluster Software Installation Guide for Solaris OS*, the command `clresource disable hadb-resource` in Step 6 is not yet available before you upgrade to Sun Cluster 3.2 software. Instead, use `scswitch -n -j hadb-resource`.

Incorrect Command to Upgrade Data Services in Live Upgrade

In “How to Upgrade the Solaris OS and Sun Cluster 3.2 Software (Live Upgrade)” in *Sun Cluster Software Installation Guide for Solaris OS*, the command in Step 19 is incorrect. To upgrade data services in a live upgrade, use the following command instead. Note that you run the `scinstall` command from the upgraded alternate BE, *not* from the installation media:

```
# BE-mount-point/usr/cluster/bin/scinstall -u update -s all \  
-d /cdrom/cdrom0/Solaris_arch/Product/sun_cluster_agents -R BE-mount-point
```

Missing Instruction to Set `kmdb` as Default x86 Boot File

In “How to Install Solaris Software” in *Sun Cluster Software Installation Guide for Solaris OS* and “How to Install Solaris and Sun Cluster Software (JumpStart)” in *Sun Cluster Software Installation Guide for Solaris OS*, the instruction to set the default x86 boot file to `kadb` applies only the Solaris 9 OS. For the Solaris 10 OS, set the default x86 boot file to `kmdb` by setting `kernel /platform/x86pc/multiboot kmdb` in the GRUB boot menu.

Additional Recovery Procedure for Dual-Partition Upgrades

You can perform the one of the following procedures instead of “How to Recover from a Failed Dual-Partition Upgrade” in *Sun Cluster Software Installation Guide for Solaris OS*. Perform these procedures if a dual-partition upgrade fails and the state of the cluster meets all of the following criteria:

- The nodes of the first partition have been upgraded.
- None of the nodes of the second partition are yet upgraded.
- None of the nodes of the second partition are in cluster mode.

You can also perform one of these procedures if the upgrade has succeeded on the first partition but you want to back out of the upgrade.

Do not perform the procedure after dual-partition upgrade processes have begun on the second partition.

Choose the procedure based on the system platform on which your cluster is installed:

SPARC Based System [“How to Recover From a Partially Completed Dual-Partition Upgrade on a SPARC Based System” on page 47](#)

x86 Based System [“How to Recover From a Partially Completed Dual-Partition Upgrade on an x86 Based System” on page 48](#)

▼ **How to Recover From a Partially Completed Dual-Partition Upgrade on a SPARC Based System**

Before You Begin Before you begin, ensure that all second-partition nodes are halted. First-partition nodes can be either halted or running in noncluster mode.

Perform all steps as superuser.

- 1 **Boot each node in the second partition into noncluster mode.**

```
# ok boot -x
```

- 2 **On each node in the second partition, run the `scinstall -u recover` command.**

```
# /usr/cluster/bin/scinstall -u recover
```

The command restores the original CCR information, restores the original `/etc/vfstab` file, and eliminates modifications for startup.

3 Boot each node of the second partition into cluster mode.

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

When the nodes of the second partition come up, the second partition resumes supporting cluster data services while running the old software with the original configuration.

4 Restore the original software and configuration data from backup media to the nodes in the first partition.

5 Boot each node in the first partition into cluster mode.

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

The nodes rejoin the cluster.

▼ How to Recover From a Partially Completed Dual-Partition Upgrade on an x86 Based System

Before You Begin Before you begin, ensure that all second-partition nodes are halted. First-partition nodes can be either halted or running in noncluster mode.

Perform all steps as superuser.

1 Boot each node in the second partition into noncluster mode by completing the following steps.

2 In the GRUB menu, use the arrow keys to select the appropriate Solaris entry and type `e` to edit its commands.

The GRUB menu appears similar to the following:

```
GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
```

```
+-----+
| Solaris 10 /sol_10_x86                |
| Solaris failsafe                      |
|                                       |
+-----+
```

Use the `^` and `v` keys to select which entry is highlighted.

Press `enter` to boot the selected OS, `'e'` to edit the commands before booting, or `'c'` for a command-line.

For more information about GRUB-based booting, see Chapter 11, “GRUB Based Booting (Tasks),” in *System Administration Guide: Basic Administration*.

- 3 In the boot parameters screen, use the arrow keys to select the kernel entry and type `e` to edit the entry.

The GRUB boot parameters screen appears similar to the following:

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
+-----+
| root (hd0,0,a)                                |
| kernel /platform/i86pc/multiboot             |
| module /platform/i86pc/boot_archive          |
+-----+
```

Use the `^` and `v` keys to select which entry is highlighted. Press `'b'` to boot, `'e'` to edit the selected command in the boot sequence, `'c'` for a command-line, `'o'` to open a new line after (`'O'` for before) the selected line, `'d'` to remove the selected line, or escape to go back to the main menu.

- 4 Add the `-x` option to the command to specify that the system boot into noncluster mode.

Minimal BASH-like line editing is supported. For the first word, TAB lists possible command completions. Anywhere else TAB lists the possible completions of a device/filename. ESC at any time exits.

```
# grub edit> kernel /platform/i86pc/multiboot -x
```

- 5 Press Enter to accept the change and return to the boot parameters screen.

The screen displays the edited command.

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
+-----+
| root (hd0,0,a)                                |
| kernel /platform/i86pc/multiboot -x         |
| module /platform/i86pc/boot_archive          |
+-----+
```

Use the `^` and `v` keys to select which entry is highlighted. Press `'b'` to boot, `'e'` to edit the selected command in the boot sequence, `'c'` for a command-line, `'o'` to open a new line after (`'O'` for before) the selected line, `'d'` to remove the selected line, or escape to go back to the main menu.-

- 6 Type `b` to boot the node into noncluster mode.

Note – This change to the kernel boot parameter command does not persist over the system boot. The next time you reboot the node, it will boot into cluster mode. To boot into noncluster mode instead, perform these steps to again to add the `-x` option to the kernel boot parameter command.

- 7 **On each node in the second partition, run the `scinstall -u recover` command.**

```
# /usr/cluster/bin/scinstall -u recover
```

The command restores the original CCR information, restores the original `/etc/vfstab` file, and eliminates modifications for startup.

- 8 **Boot each node of the second partition into cluster mode.**

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

When the nodes of the second partition come up, the second partition resumes supporting cluster data services while running the old software with the original configuration.

- 9 **Restore the original software and configuration data from backup media to the nodes in the first partition.**

- 10 **Boot each node in the first partition into cluster mode.**

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

The nodes rejoin the cluster.

Sun Cluster Data Services Planning and Administration Guide

Support of Scalable Services on Non-Global Zones

In “Resource Type Properties” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*, the description of the `Failover` resource property is missing a statement concerning support of scalable services on non-global zones. This support applies to resources for which the `Failover` property of the resource type is set to `FALSE` and the `Scalable` property of the resource is set to `TRUE`. This combination of property settings indicates a scalable service that uses a `SharedAddress` resource to do network load balancing. In the Sun Cluster 3.2 release, you can configure a scalable service of this type in a resource group that runs in a non-global zone. But you cannot configure a scalable service to run in multiple non-global zones on the same node.

System Administration Guide

This section discusses errors and omissions in the *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

Incorrect Command to Shut Down Cluster Nodes

In Example 3-5 in “How to Shut Down a Cluster Node” in *Sun Cluster System Administration Guide for Solaris OS*, the use of the option `-T0` with the shutdown command is incorrect. The correct option is `-g0`.

Incorrect Step in How to Reboot a Cluster Node

The procedure “How to Back Up the Root (/) File System” in *Sun Cluster System Administration Guide for Solaris OS* provides the incorrect instruction to boot into noncluster mode on an x86 based system. Instead, follow instructions for the GRUB based booting method in “How to Boot a Cluster Node in Noncluster Mode” in *Sun Cluster System Administration Guide for Solaris OS*.

Incorrect Instruction to Boot an x86 Based System Into Noncluster Mode

The procedure “How to Back Up the Root (/) File System” in *Sun Cluster System Administration Guide for Solaris OS* provides the incorrect instruction to boot into noncluster mode on an x86 based system. Instead, follow instructions for the GRUB based booting method in “How to Boot a Cluster Node in Noncluster Mode” in *Sun Cluster System Administration Guide for Solaris OS*.

Incorrect Instruction to Reboot an x86 Based System Into Cluster Mode

The procedure “How to Reboot a Cluster Node” in *Sun Cluster System Administration Guide for Solaris OS* provides incorrect information. See the following list for the correct information.

Step 3 The steps provides the wrong command. Use the shutdown `-g0 -y -i6` command.

Example 3-9 The example shows the following incorrect instructions to reboot a cluster node in an x86 based system. Instead, follow instructions for the GRUB based booting method in “How to Boot a Cluster Node” in *Sun Cluster System Administration Guide for Solaris OS*.

Extraneous Step In Procedure to Remove Connectivity

The procedure “How to Remove Connectivity Between an Array and a Single Node, in a Cluster With Greater Than Two-Node Connectivity” in *Sun Cluster System Administration Guide for Solaris OS* contains an unneeded step. Ignore Step 11, which instructs you to boot into noncluster mode and contains obsolete instructions for x86 based systems. Step 13 contains the correct boot instructions.

Sun Cluster Module to Sun Management Center No Longer Provides Online Help

The sections “Administration Tools” in *Sun Cluster System Administration Guide for Solaris OS*, “Beginning to Administer the Cluster” in *Sun Cluster System Administration Guide for Solaris OS*, and “Overview of Sun Management Center” in *Sun Cluster System Administration Guide for Solaris OS* contain references to online help for the Sun Cluster module to Sun Management Center. This information is no longer valid for the Sun Cluster 3.2 release. The version of Sun Management Center that Sun Cluster 3.2 software supports no longer provides or supports online help. For information about how to navigate within Sun Management Center windows, see the Sun Management Center documentation. For definitions of cluster objects that the Sun Cluster module monitors, see the *Sun Cluster Overview for Solaris OS* or *Sun Cluster Concepts Guide for Solaris OS*.

Location of SNMP MIB Definition File

The SNMP event MIB is defined in the `sun-cluster-event-mib.mib` file and is located in the `/usr/cluster/lib/mib` directory. You can use this definition to interpret the SNMP trap information.

Data Services Developer's Guide

This section discusses errors and omissions in the *Sun Cluster Data Services Developer's Guide for Solaris OS*.

Support of Certain Scalable Services on Non-Global Zones

In “Resource Type Properties” in *Sun Cluster Data Services Developer's Guide for Solaris OS*, the description of the `FaiLover` resource property is missing a statement concerning support of scalable services on non-global zones. This support applies to resources for which the `FaiLover` property of the resource type is set to `FALSE` and the `Scalable` property of the resource is set to `TRUE`. This combination of property settings indicates a scalable service that uses a `SharedAddress` resource to do network load balancing. In the Sun Cluster 3.2 release, you can configure a scalable service of this type in a resource group that runs in a non-global zone. But you cannot configure a scalable service to run in multiple non-global zones on the same node.

Method Timeout Behavior is Changed

A description of the change in the behavior of method timeouts in the Sun Cluster 3.2 release is missing. If an RGM method callback times out, the process is now killed by using the `SIGABRT` signal instead of the `SIGTERM` signal. This causes all members of the process group to generate a core file.

Note – Avoid writing a data-service method that creates a new process group. If your data service method does need to create a new process group, also write a signal handler for the SIGTERM and SIGABRT signals. Write the signal handlers to forward the SIGTERM or SIGABRT signal to the child process group before the signal handler terminates the parent process. This increases the likelihood that all processes that are spawned by the method are properly terminated.

CRNP Runs Only in the Global Zone

Chapter 12, “Cluster Reconfiguration Notification Protocol,” in *Sun Cluster Data Services Developer’s Guide for Solaris OS* is missing the statement that, on the Solaris 10 OS, the Cluster Recognition Notification Protocol (CRNP) runs only in the global zone.

Required Solaris Software Group Statement is Unclear

In “Setting Up the Development Environment for Writing a Data Service” in *Sun Cluster Data Services Developer’s Guide for Solaris OS*, there is a Note that the Solaris software group Developer or Entire Distribution is required. This statement applies to the development machine. But because it is positioned after a statement about testing the data service on a cluster, it might be misread as a requirement for the cluster that the data service is being run on.

Data Service for SAP Guide

This section discusses errors and omissions in the *Sun Cluster Data Service for SAP Guide for Solaris OS*.

Changes to How to Install and Enable an SAP Scalable Application Server

The following procedure has been corrected.

▼ How to Install and Enable an SAP Scalable Application Server

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 the 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” in *Sun Cluster Data Service for SAP Guide for Solaris OS*.

- 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 about 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/SUN/D03

# chown sapsidadm:sapsys /usr/sap/local/SUN/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. The two application servers will then 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/SAPSID/D03 /usr/sap/SAPSID/D03
```

5 Install the application server on all nodes.

6 Are you using SAP 6.10, SAP 6.20, or SAP 6.40?

- 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 file names that you choose specify this instance.

```
# cp /usr/sap/SAPSID/SYS/exe/run/startsap \
  $SAPSID_HOME/startsap_instance-number

# cp /usr/sap/SAPSID/SYS/exe/run/stopsap \
  $SAPSID_HOME/stopsap_instance-number
```

- 9 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* or *oradbsid*.**

Example:

```
# vi startsap_03
```

Before:

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

After:

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

- 10 Are you using the MaxDB database?**

- If no, skip to Step 17.
- If yes, proceed to Step 15.

- 11 Create a `.XUSER.62` file in the *SAPSIDADm* home directory as described in “How to Install and Configure MaxDB” in *Sun Cluster Data Service for MaxDB Guide for Solaris OS*.**

Note – This user whose home directory contains the `.XUSER.62` file is an internal database user who has permissions to start, stop, and query the database.

- 12 Verify that the `.XUSER.62` file is correctly configured by using the `-d` option of the `R3trans` command. You must be an *SAPSIDADm* user to perform this verification.**

- 13 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 resource group `nodeList`, which is located in the scalable application server.

- 14 Ensure that you can start and stop the application server from each node, and verify that the log files are in the correct location. For further information about updates on SAP profiles, see the SAP documentation at <http://service.sap.com/ha>.**

Note – Update the files under the *SAPSIDADm* home directory so that they refer to the corresponding logical hosts. This ensures that the necessary environment variables are sourced.

- 15 Create the SAP logon group if you use a logon group.**

Quorum Server User's Guide

This section discusses errors and omissions in the *Sun Cluster Quorum Server User's Guide*.

Supported Software and Hardware Platforms

The following installation requirements and guidelines are missing or unclear:

- Solaris software requirements for Sun Cluster software apply as well to Quorum Server software.
- The supported hardware platforms for a quorum server are the same as for a cluster node.
- A quorum server does not have to be configured on the same hardware and software platform as the cluster or clusters that it provides quorum to. For example, an x86 based machine that runs the Solaris 9 OS can be configured as a quorum server for a SPARC based cluster that runs the Solaris 10 OS.
- A quorum server can be configured on a cluster node to provide quorum for clusters other than the cluster that the node belongs to. However, a quorum server that is configured on a cluster node is not highly available.

Man Pages

This section discusses errors and omissions in the Sun Cluster man pages.

`clnode(1CL)`

- The description of the `remove` subcommand implies that the command will not work when certain conditions exist. Instead, the command will execute in these conditions but the results might adversely affect the cluster. The following is a more accurate description of the `remove` subcommand requirements and behavior:

To remove a node from a cluster, observe the following guidelines. If you do not observe these guidelines, the removal of a node might compromise quorum in the cluster.

- Unconfigure the node to be removed from any quorum devices, unless you also specify the `-f` option.
- Ensure that the node to be removed is not an active cluster member.
- Do not remove a node from a three-node cluster unless at least one shared quorum device is configured.

The `clnode remove` command attempts to remove a subset of references to the node from the cluster configuration database. If the `-f` option is also specified, the subcommand attempts to remove all references to the node.

- Before you can successfully use the `clnode remove` command to remove a node from the cluster, you must first use the `claccess add` command to add the node to the cluster authentication list, if it is not already in the list. Use the `claccess list` or `claccess show` command to view the current cluster authentication list. Afterwards, for security use the `claccess deny -all` command to prevent further access of the cluster configuration by any cluster node. For more information, see the `claccess(1CL)` man page.

clresource(1CL)

- The following option is missing from the `clresource(1CL)` man page:

`-u` Specifies that the command operate on resources whose resource group is suspended, if you specify the `+` operand. If you do not also specify the `u` option when you specify the `+` operand, the command ignores all resources whose resource group is suspended.

The `-u` option is valid when the `+` operand is specified to the `clear`, `disable`, `enable`, `monitor`, `set`, and `unmonitor` subcommands.

- The description of the `+` operand should state that, when used with the `clear`, `disable`, `enable`, `monitor`, `set`, or `unmonitor` subcommand, the command ignores all resources whose resource group is suspended, unless you also specify the `-u` option.
- The example provided in the definitions of the `+` and `-` operands for the `-p`, `-x`, and `-y` options are incorrect. The definitions should be as follows:
 - `+` Adds a value or values to a string array value. Only the `set` subcommand accepts this operator. You can specify this operator only for the properties that accept lists of string values, for example `Resource_dependencies`.
 - `-` Deletes a value or values from a string array value. Only the `set` subcommand accepts this operator. You can specify this operator only for properties that accept lists of string values, for example `Resource_dependencies`.

clresourcegroup(1CL)

- The command syntax and description for the `evacuate` subcommand incorrectly states that you can evacuate more than one node or zone in the same command invocation. Instead, you can specify only one node or zone in the `evacuate` command
- The following option is missing from the `clresourcegroup(1CL)` man page:
 - `-u` Specifies that the command operate on suspended resource groups, if you specify the `+` operand. If you do not also specify the `u` option when you specify the `+` operand, the command ignores all suspended resource groups.

The `-u` option is valid when the `+` operand is specified to the `add-node`, `manage`, `offline`, `online`, `quiesce`, `remaster`, `remove-node`, `restart`, `set`, `switch`, and `unmanage` subcommands.

- The description of the `+` operand should state that, when used with the `add-node`, `manage`, `offline`, `online`, `quiesce`, `remaster`, `remove-node`, `restart`, `set`, `switch`, or `unmanage` subcommand, the command ignores all suspended resource groups, unless you also specify the `-u` option.

r_properties(5)

- The use of the `Network_resources_used` property has changed in the Sun Cluster 3.2 release. If you do not assign a value to this property, its value is updated automatically by the RGM, based on the setting of the resource-dependencies properties. You do not need to set this property directly. Instead, set the `Resource_dependencies`, `Resource_dependencies_offline_restart`, `Resource_dependencies_restart`, or `Resource_dependencies_weak` property.

To maintain compatibility with earlier releases of Sun Cluster software, you can still set the value of the `Network_resources_used` property directly. If you do, the value of the `Network_resources_used` property is no longer derived from the settings of the resource-dependencies properties.

If you add a resource name to the `Network_resources_used` property, the resource name is automatically added to the `Resource_dependencies` property as well. The only way to remove that dependency is to remove it from the `Network_resources_used` property. If you are not sure whether a network-resource dependency was originally added to the `Resource_dependencies` property or to the `Network_resources_used` property, remove the dependency from both properties. For example, the following command removes a dependency of resource `r1` upon network resource `r2`, regardless of whether the dependency was added to the `Network_resources_used` property or to the `Resource_dependencies` property:

```
# clresource set -p Network_resources_used-=r2 -p Resource_dependencies-=r2 r1
```

- The `r_properties(5)` man page contains incorrect descriptions of the `Resource_dependencies`, `Resource_dependencies_offline_restart`, `Resource_dependencies_restart`, and `Resource_dependencies_weak` properties. For correct descriptions of these properties, instead see “Resource Properties” in *Sun Cluster Data Services Developer’s Guide for Solaris OS*.
- The description of the `Scalable` resource property is missing a statement concerning support of scalable services on non-global zones. This support applies to resources for which the `Failover` property of the resource type is set to `FALSE` and the `Scalable` property of the resource is set to `TRUE`. This combination of property settings indicates a scalable service that uses a `SharedAddress` resource to do network load balancing. In the Sun Cluster 3.2 release,

you can configure a scalable service of this type in a resource group that runs in a non-global zone. But you cannot configure a scalable service to run in multiple non-global zones on the same node.

rt_properties(5)

The description of the `Failover` resource-type property contains an incorrect statement concerning support of scalable services on non-global zones in the Sun Cluster 3.2 release. This applies to resources for which the `Failover` property of the resource type is set to `FALSE` and the `Scalable` property of the resource is set to `TRUE`.

Incorrect: You cannot use a scalable service of this type in zones.

Correct: You can configure a scalable service of this type in a resource group that runs in a non-global zone. But you cannot configure a scalable service to run in multiple non-global zones on the same node.

SUNW.Event(5)

The `SUNW.Event(5)` man page is missing the statement that, on the Solaris 10 OS, the Cluster Recognition Notification Protocol (CRNP) runs only in the global zone.

