



Sun Cluster 3.2 Release Notes for Solaris OS



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

Part No: 819-6611-12
May 2007, Revision A

Copyright 2007 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 2007 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 dérivés 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 CONTREFACON.

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
Restrictions	17
Compatibility Issues	17
Features Nearing End of Life	17
Solaris 10 11/06 Operating System	18
fssnap Support	19
Solaris Volume Manager GUI	19
Loopback File System (LOFS)	19
Accessibility Features for People With Disabilities	19
Commands Modified in This Release	19
Object-Oriented Command Line Interface	20
scinstall Command	20
scconf Command	20
Resource Properties	21
Product Name Changes	21
Supported Products	21
Sun Cluster Security Hardening	29
Known Issues and Bugs	30
Administration	30
Data Services	33
Installation	33
Localization	35
Runtime	35
Upgrade	36

Patches and Required Firmware Levels	41
Applying the Sun Cluster 3.2 Core Patch	42
▼ How to Apply the Sun Cluster 3.2 Core Patch	42
Removing the Sun Cluster 3.2 Core Patch	42
▼ How to Remove the Sun Cluster 3.2 Core Patch	43
Patch Management Tools	43
SunSolve Online	44
Sun Cluster 3.2 Documentation	45
Searching Sun Product Documentation	45
Sun Cluster 3.2 Software Manuals for Solaris OS	45
Sun Cluster 3.2 Data Service Manuals for Solaris OS (SPARC Platform Edition)	46
Sun Cluster 3.2 Data Service Manuals for Solaris OS (x86 Platform Edition)	48
Sun Cluster 3.1 - 3.2 Hardware Collection for Solaris OS (SPARC Platform Edition)	49
Sun Cluster 3.1 — 3.2 Hardware Collection for Solaris OS (x86 Platform Edition)	50
Documentation Issues	50
Concepts Guide	50
Software Installation Guide	51
Sun Cluster Data Services Planning and Administration Guide	51
Sun Cluster Data Service for MaxDB Guide	52
Sun Cluster Data Service for SAP Guide	52
Sun Cluster Data Service for SAP liveCache Guide	53
Sun Cluster Data Service for SAP Web Application Server Guide	54
System Administration Guide	58
Data Services Developer's Guide	61
Quorum Server User's Guide	62
Man Pages	62

Preface

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% you have mail.</code>
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name% su</code> Password:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <code>rm filename</code> .
<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 17
- “Commands Modified in This Release” on page 19
- “Product Name Changes” on page 21
- “Supported Products” on page 21
- “Known Issues and Bugs” on page 30
- “Patches and Required Firmware Levels” on page 41
- “Sun Cluster 3.2 Documentation” on page 45
- “Documentation Issues” on page 50

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 11
- “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 12
- “Sun Cluster Quorum Server” on page 12

- “Disk-Path Failure Handling” on page 12
- “HASToragePlus Availability Improvements” on page 12
- “Solaris Zones Expanded Support” on page 12
- “ZFS” on page 13
- “HDS TrueCopy Campus Cluster” on page 13
- “Specifications-Based Campus Cluster” on page 13
- “Multi-Terabyte Disk and Extensible Firmware Interface (EFI) Label Support” on page 14
- “Extended Support for VERITAS Software Components” on page 14
- “Quota Support” on page 14
- “Oracle DataGuard Support” on page 14
- “Dual-Partition Upgrade ” on page 14
- “Live Upgrade” on page 15
- “Optional Sun Cluster Manager Installation” on page 15
- “SNMP Event MIB” on page 15
- “Command Logging” on page 15
- “Workload System-Resource Monitoring” on page 15
- “Automatic Creation of Multiple-Adapter IPMP Groups by `scinstall`” on page 16
- “Secure Shell Support for Cluster Control Panel Software” on page 16
- “New Minimum Requirement of One Cluster Interconnect” on page 16
- “IP Filter Support for Failover Services ” on page 16

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
- Oracle RAC
- HA-NFS
- HA-Apache, all versions shipped with Solaris software
- HA-SAP

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

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* 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* 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 the 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 paths 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* for more information.

HASStoragePlus Availability Improvements

HASStoragePlus 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 support 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 the Sun Cluster 3.2 release. 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.

Specifications-Based Campus Cluster

Specifications-Based Campus Clusters now support a wider range of distance configurations. These clusters support such configurations by requiring compliance to a latency and error rate, rather than to a rigid set of distances and components.

See Chapter 7, “Campus Clustering With Sun Cluster Software,” in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS* for more information.

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 new feature 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

Sun Cluster software now offers improved usability for Oracle deployments including DataGuard data replication software. Customers can now specify an HA-Oracle database to be part of an Oracle DataGuard configuration as either a primary or a standby site. This secondary database can be a logical or a physical standby. For more information, see *Sun Cluster Data Service for Oracle Guide for Solaris OS*.

Note – When the HA-Oracle agent is managing a standby database, the agent will only control start, stop, and monitoring of that database. The agent does not re-initiate the recovery of the standby database if it fails over to another node.

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 outages 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 36 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* 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. These new features 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*.

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 now configures either a single-adapter or a multiple-adapter IPMP group for each set of public-network adapters depending on the adapters available in each subnet. This functionality replaces the utility's previous behavior which created one single-adapter IPMP group for each adapter available regardless of their subnet. For more information about this and other changes to IPMP group policies, see “Public Networks” in *Sun Cluster Software Installation Guide for Solaris OS*.

Secure Shell Support for Cluster Control Panel Software

Support for Secure Shell is added to the Cluster Control Panel (CCP) by the following new features:

- Addition of Secure Shell support to the `cconsole` utility. To make Secure Shell connections to node consoles from the `cconsole` graphical user interface (GUI), enable the Use SSH checkbox in the Options menu.

Alternatively, you can launch the utility in Secure Shell mode directly, by typing the following command from the command line:

```
cconsole -s [-l username]
```

- Introduction of the new `cssh` utility, to connect securely to cluster nodes.

For more information about preparing for and using the Secure Shell features of the CCP, see “How to Install Cluster Control Panel Software on an Administrative Console” in *Sun Cluster Software Installation Guide for Solaris OS*. For updates to the related man pages, see “`ccp(1M)`” on page 62, “`cconsole(1M)`, `crlogin(1M)`, `cssh(1M)`, and `ctelnet(1M)`” on page 63, and “`serialports(4)`” on page 67.

New Minimum Requirement of One Cluster Interconnect

The minimum required number of cluster interconnects that a cluster must have is changed to one cluster interconnect between the nodes of the cluster. The interactive `scinstall` utility is revised to permit configuration of only one interconnect when you use the utility in Custom mode. To use the utility's Typical mode, you must still configure two interconnects. For more information, see “Cluster Interconnect” in *Sun Cluster Software Installation Guide for Solaris OS*.

IP Filter Support for Failover Services

Sun Cluster 3.2 software supports the Solaris IP Filter for failover services. Solaris IP Filter provides stateful packet filtering and network address translation (NAT). Solaris IP Filter also includes the ability to create and manage address pools. For more information on the Solaris IP Filter, see Part IV, “IP Security,” in *System Administration Guide: IP Services*. For information on how to set up IP filtering with Sun Cluster software, see “Using Solaris IP Filtering with Sun Cluster” on page 60.

Restrictions

NetApp NAS Fencing Restriction

The fencing feature requires that each cluster node always use the same source IP address when accessing the NetApp NAS unit. Multi-homed systems use multiple source IP addresses. The administrator for a multi-homed system must ensure that one source IP address is always used when accessing the NetApp NAS unit. This can be achieved by setting up an appropriate network configuration.

Compatibility Issues

- “Features Nearing End of Life” on page 17
- “Solaris 10 11/06 Operating System” on page 18
- “fssnap Support” on page 19
- “Solaris Volume Manager GUI” on page 19
- “Loopback File System (LOFS)” on page 19
- “Accessibility Features for People With Disabilities” on page 19

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 30.

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 discontinued. The Sun Cluster 3.0 part number will no longer be available.

Solaris 8

As of Sun Cluster 3.2, Sun Cluster will no 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.

sccheck

The `sccheck` command might not be included in a future release. However, the corresponding functionality will be provided by the `cluster check` command.

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 Solaris patches to fix these issues. For more information, refer to Infodoc 87995.



Caution – You must upgrade your operating system to Solaris 10 11/06 *before* applying the Solaris patches.

- 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.

fssnap Support

Sun Cluster does not support fssnap which is a feature of UFS. You can use fssnap on local systems that are not controlled by Sun Cluster. The following restrictions apply to fssnap support:

- Supported on local filesystems not managed by Sun Cluster software
- Not supported on global filesystems
- Not supported on local filesystems under the control of HAStoragePlus

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.

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 41](#) 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 “[r_properties\(5\)](#)” on page 66.

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.

Note – Sun Cluster 3.2 software is distributed under Solaris Cluster 3.2 and Sun Java Availability Suite.

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 that are 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 MP2.	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 MP2.	VERITAS Volume Manager 4.1 cluster feature.
	Solaris 10	VERITAS Volume Manager components that are delivered as part of VERITAS Storage Foundation 5.0. This support requires VxVM 5.0 MP1.	VERITAS Volume Manager 5.0 cluster feature.
		Solaris Volume Manager.	Solaris Volume Manager for Sun Cluster.
		VERITAS Volume Manager 4.1. This support requires VxVM 4.1 MP2.	VERITAS Volume Manager 4.1 with cluster feature.
	Solaris 10	VERITAS Volume Manager 4.1. This support requires VxVM 4.1 MP2.	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 VxVM 5.0 MP1.	VERITAS Volume Manager 5.0 cluster feature.
		Solaris Volume Manager.	Solaris Volume Manager for Sun Cluster.
x86	Solaris 10	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.
		Solaris Volume Manager.	Solaris Volume Manager for Sun Cluster.

■ File systems

Platform	Operating System	File System	Features and External Volume Management
SPARC	Solaris 9	Solaris UFS.	N/A
		Sun StorEdge QFS:	N/A
		QFS 4.5 Standalone Filesystem.	Features: <ul style="list-style-type: none">■ HA-NFS■ HA-Oracle External Volume Management: <ul style="list-style-type: none">■ SVM■ VxVM
		QFS 4.5 — Shared QFS File System.	Feature: <ul style="list-style-type: none">■ Oracle RAC External Volume Management: <ul style="list-style-type: none">■ SVM Cluster File Manager
		QFS 4.6.	Features: <ul style="list-style-type: none">■ COTC-Shared QFS clients outside the cluster■ HA-SAM Failover
		VERITAS File System 4.1.	N/A
		VERITAS File System components that are delivered as part of VERITAS Storage Foundation 4.1 and 5.0.	N/A

Platform	Operating System	File System	Features and External Volume Management
SPARC	Solaris 10	Solaris ZFS.	N/A
		Solaris ZFS.	N/A
		Sun StorEdge QFS:	N/A
		QFS 4.5 Standalone Filesystem.	Features: <ul style="list-style-type: none"> ■ HA-NFS ■ HA-Oracle External Volume Management: <ul style="list-style-type: none"> ■ SVM ■ VxVM
		QFS 4.5 — Shared QFS File System.	Feature: <ul style="list-style-type: none"> ■ Oracle RAC External Volume Management: <ul style="list-style-type: none"> ■ SVM Cluster File Manager
		QFS 4.6.	Features: <ul style="list-style-type: none"> ■ COTC-Shared QFS clients outside the cluster ■ HA-SAM Failover
		VERITAS File System 4.1.	N/A
		VERITAS File System components that are delivered as part of VERITAS Storage Foundation 4.1 and 5.0.	N/A

Platform	Operating System	File System	Features and External Volume Management
x86	Solaris 10	Solaris UFS.	N/A
		Solaris ZFS.	N/A
		Sun StorEdge QFS:	N/A
		QFS 4.5 Standalone Filesystem	Features: <ul style="list-style-type: none">■ HA-NFS■ HA-Oracle External Volume Management: <ul style="list-style-type: none">■ SVM■ VxVM
		QFS 4.5 — Shared QFS File System.	Feature: <ul style="list-style-type: none">■ Oracle RAC External Volume Management: <ul style="list-style-type: none">■ SVM Cluster File Manager
QFS 4.6.	Features: <ul style="list-style-type: none">■ COTC-Shared QFS clients outside the cluster■ HA-SAM Failover		

- **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.

Note – Data service documentation, including man pages and wizard online help, is no longer translated to languages other than English.

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.gds
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

Data Service	Sun Cluster Resource Type
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
	SUNW.scaldevicegroup
	SUNW.scalmountpoint
	SUNW.crs_framework
SUNW.scalable_rac_server_proxy	
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
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

Data Service	Sun Cluster Resource Type
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 support 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 N1 GridEngine
- 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 SWIFTAlliance Gateway
- 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 and issues are grouped into the following categories:

- [“Administration” on page 30](#)
- [“Data Services” on page 33](#)
- [“Installation” on page 33](#)
- [“Localization” on page 35](#)
- [“Runtime” on page 35](#)
- [“Upgrade” on page 36](#)

Administration

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 SUNWCluster 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.

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 Start Quorum Server software.

```
# 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.

New Variable in HA MySQL Patch Must be Configured for All New Instances (6516322)

Problem Summary: The HA MySQL patch adds a new variable called `MYSQL_DATADIR` in the `mysql_config` file. This new variable must point to the directory where the MySQL configuration file `my.conf` file is stored. If this variable is not configured correctly, the database preparation with `mysql_register` will fail.

Workaround: Point the `MYSQL_DATADIR` variable to the directory where the MySQL configuration file, `my.conf` is stored.

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: 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 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 directly connected.

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 `cluster` entry 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 `cluster` entry 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/nsswitch.conf` file for non-global zones with the `cluster` entry 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

Translated Messages for Quorum Server are Delivered as Part of the Core Translation Packages (6482813)

Problem Summary: Translated messages for the Quorum Server administration programs, such as `clquorumserver`, are delivered as part of the core translation packages. As a result, Quorum Server messages appear only in English. The Quorum server translation packages must be separated from the core translation packages and installed on the quorum server system.

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 `SUNWjscman` (Japanese man page) package.

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.

Runtime

SAP `cleanipc` Binary Needs `User_env` Parameter for `LD_LIBRARY_PATH` (4996643)

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 the `ld.config` file.

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

```
# crle -u -l /sapmnt/SAPSID/exe
```

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

```
# crle -64 -u -l /sapmnt/SAPSID/exe
```

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

Upgrade

The `vxlu finish` 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 `vxlu finish` 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: ugettxt: 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.

Live Upgrade Should Support Mounting Global Devices From Boot Disk (6433728)

Problem Summary: During live upgrade, the `luc create` and `luupgrade` commands fail to change the DID names in the alternate boot environment that corresponds to the `/global/.devices/node@N` entry.

Workaround: Before you start the live upgrade, perform the following steps on each cluster node.

1. Become superuser.
2. Back up the `/etc/vfstab` file.


```
# cp /etc/vfstab /etc/vfstab.old
```
3. Open the `/etc/vfstab` file for editing.
4. Locate the line that corresponds to `/global/.device/node@N`.
5. Edit the global device entry.
 - Change the DID names to the physical names.
Change `/dev/did/{r}dsk/dYsZ` to `/dev/{r}dsk/cNtXdYsZ`.
 - Remove `global` from the entry.

The following example shows the name of DID device `d3s3` which corresponds to `/global/.devices/node@s`, changed to its physical device names and the `global` entry removed:

Original:

```
/dev/did/dsk/d3s3 /dev/did/rdisk/d3s3 /global/.devices/node@2 ufs 2 no global
```

Changed:

```
dev/dsk/c0t0d0s3 /dev/rdsk/c0t0d0s3 /global/.devices/node@2 ufs 2 no -
```

6. When the `/etc/vfstab` file is modified on all cluster nodes, perform live upgrade of the cluster, but stop before you reboot from the upgraded alternate boot environment.
7. On each node, on the current, unupgraded, boot environment, restore the original `/etc/vfstab` file.

```
# cp /etc/vstab.old /etc/vfstab
```

8. In the alternate boot environment, open the `/etc/vfstab` file for editing.
9. Locate the line that corresponds to `/global/.devices/node@N` and replace the dash (-) at the end of the entry with the word `global`.

```
/dev/dsk/cNtXdYsZ /dev/rdsk/cNtXdYsZ /global/.devices/node@N ufs 2 no global
```

10. Reboot the node from the upgraded alternate boot environment.

The DID names are substituted in the `/etc/vfstab` file automatically.

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.

Cannot Use Zones Following Live Upgrade From Sun Cluster Version 3.1 on Solaris 9 to Version 3.2 on Solaris 10 (6509958)

Problem Summary: Following a live upgrade from Sun Cluster version 3.1 on Solaris 9 to version 3.2 on Solaris 10, zones cannot be used properly with the cluster software. The problem is that the `pspool` data is not created for the Sun Cluster packages. So those packages that must be propagated to the non-global zones, such as `SUNWsczu`, are not propagated correctly.

Workaround: After the Sun Cluster packages have been upgraded by using the `scinstall -R` command but before the cluster has booted into cluster mode, run the following script twice:

- Once for the Sun Cluster framework packages
- Once for the Sun Cluster data-service packages

▼ Instructions for Using the Script

Before You Begin Prepare and run this script in one of the following ways:

- Set up the variables for the Sun Cluster framework packages and run the script. Then modify the `PATHNAME` variable for the data service packages and rerun the script.
- Create two scripts, one with variables set in the script for the framework packages and one with variables set for the data service packages. Then run both scripts.

1 Become superuser.

2 Create a script with the following content.

```
#!/bin/ksh

typeset PLATFORM=${PLATFORM:-'uname -p'}
typeset PATHNAME=${PATHNAME:-/cdrom/cdrom0/Solaris_${PLATFORM}/Product/sun_cluster/Solaris_10/Packages}
typeset BASEDIR=${BASEDIR:-/}

cd $PATHNAME
for i in *
do
    if pkginfo -R ${BASEDIR} $i >/dev/null 2>&1
    then
        mkdir -p ${BASEDIR}/var/sadm/pkg/$i/save/pspool
        pkgadd -d . -R ${BASEDIR} -s ${BASEDIR}/var/sadm/pkg/$i/save/pspool $i
    fi
done
```

3 Set the variables PLATFORM, PATHNAME, and BASEDIR.

Either set these variables as environment variables or modify the values in the script directly.

PLATFORM The name of the platform. For example, it could be `sparc` or `x86`. By default, the `PLATFORM` variable is set to the output of the `uname -p` command.

PATHNAME A path to the device from where the Sun Cluster framework or data-service packages can be installed. This value corresponds to the `-d` option in the `pkgadd` command.

As an example, for Sun Cluster framework packages, this value would be of the following form:

```
/cdrom/cdrom0/Solaris_${PLATFORM}/Product/sun_cluster/Solaris_10/Packages
```

For the data services packages, this value would be of the following form:

```
/cdrom/cdrom0/Solaris_${PLATFORM}/Product/sun_cluster_agents/Solaris_10/Packages
```

BASEDIR The full path name of a directory to use as the root path and corresponds to the `-R` option in the `pkgadd` command. For live upgrade, set this value to the root path that is used with the `-R` option in the `scinstall` command. By default, the `BASEDIR` variable is set to the root (`/`) file system.

4 Run the script, once for the Sun Cluster framework packages and once for the data-service packages.

After the script is run, you should see the following message at the command prompt for each package:

```
Transferring pkgname package instance
```

Note – If the `pspool` directory already exists for a package or if the script is run twice for the same set of packages, the following error is displayed at the command prompt:

```
Transferring pkgname package instance
pkgadd: ERROR: unable to complete package transfer
  - identical version of pkgname already exists on destination device
```

This is a harmless message and can be safely ignored.

- 5 **After you run the script for both framework packages and data-service packages, boot your nodes into cluster mode.**

Can't Add Node to an Existing Sun Cluster 3.2–Patched Cluster Without Adding the Sun Cluster 3.2 Core Patch to the Node (6554107)

Problem Summary: Adding a new cluster node without ensuring that the node has the same patches as the existing cluster nodes might cause the cluster nodes to panic.

Workaround: Before adding nodes to the cluster, ensure that the new node is first patched to the same level as the existing cluster nodes. Failure to do this might cause the cluster nodes to panic.

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 – Read the patch README before applying or removing any patch.

If you are using the rebooting patch (node) method to install the Sun Cluster core patch, 125510 (S9/SPARC), 125511 (S10/SPARC), or 125512 (S19/x64), you must have the -02 version of the patch installed before you can install higher versions of the patch. If you do not have the -02 patch installed and wish to install -03 or higher (when available) you must use the rebooting cluster method.

See the following list for examples of patching scenarios:

- If you have Sun Cluster 3.2 software using the Solaris 10 operating system on SPARC with patch 125511-02 and wish to install 125511-03 or higher, you may use the rebooting node or rebooting cluster method.
- If you have Sun Cluster 3.2 software using the Solaris 10 operating system on SPARC without 125511-02 installed and wish to install 125511-03 or higher, the choices are:
 - Use the rebooting cluster method to install 125511-03.
 - Install 125511-02 using the rebooting node method and then install 125511-03 using the rebooting node method.

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>.

Applying the Sun Cluster 3.2 Core Patch

Complete the following procedure to apply the Sun Cluster 3.2 core patch.

▼ How to Apply the Sun Cluster 3.2 Core Patch

- 1 Install the patch using the usual rebooting patch procedure for a core patch.
- 2 Verify that the patch has been installed correctly on all nodes and is functioning properly.
- 3 **Register the new version of resource types `SUNW.HASStoragePlus`, `SUNW.ScalDeviceGroup`, and `SUNW.ScalMountPoint` that are being updated in this patch. Perform resource type upgrade on any existing resources of these types to the new versions.**

For information about registering a resource type, see “Registering a Resource Type” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.



Caution – If the Sun Cluster 3.2 core patch is removed, any resources that were upgraded in step 3 must be downgraded to the earlier resource type versions. The procedure for downgrading will require planned downtime of these services. Therefore, do not perform step 3 until you are ready to commit the Sun Cluster 3.2 core patch permanently to your cluster.

Removing the Sun Cluster 3.2 Core Patch

Complete the following procedure to remove the Sun Cluster 3.2 core patch.

▼ How to Remove the Sun Cluster 3.2 Core Patch

- 1 List the resource types on the cluster.

```
# clrt list
```

- 2 If the list returns `SUNW.HASStoragePlus:5`, `SUNW.ScalDeviceGroup:2`, or `SUNW.ScalMountPoint:2`, you must remove these resource types. For instructions on removing a resource type, see “How to Remove a Resource Type” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

- 3 Reboot all nodes of the cluster into noncluster, single user mode.

For instructions on rebooting cluster nodes into noncluster, single user mode, see “How to Boot a Cluster Node in Noncluster Mode” in *Sun Cluster System Administration Guide for Solaris OS* in *Sun Cluster System Administration Guide for Solaris OS*.

- 4 Remove the Sun Cluster 3.2 core patch from each node on which you installed the patch.

```
# patchrm patch-id
```

- 5 Reboot into cluster mode all of the nodes from which you removed the Sun Cluster 3.2 core patch.

Rebooting all of the nodes from which you removed the Sun Cluster 3.2 core patch before rebooting any unaffected nodes ensures that the cluster is formed with the correct information in the CCR. If all nodes on the cluster were patched with the core patch, you can reboot the nodes into cluster mode in any order.

For instructions on rebooting nodes into cluster mode, see “How to Reboot a Cluster Node” in *Sun Cluster System Administration Guide for Solaris OS* in *Sun Cluster System Administration Guide for Solaris OS*.

- 6 Reboot any remaining nodes into cluster mode.

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://ttp://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:

- “Sun Cluster 3.2 Software Manuals for Solaris OS” on page 45
- “Sun Cluster 3.2 Data Service Manuals for Solaris OS (SPARC Platform Edition)” on page 46
- “Sun Cluster 3.2 Data Service Manuals for Solaris OS (x86 Platform Edition)” on page 48
- “Sun Cluster 3.1 - 3.2 Hardware Collection for Solaris OS (SPARC Platform Edition)” on page 49
- “Sun Cluster 3.1 — 3.2 Hardware Collection for Solaris OS (x86 Platform Edition)” on page 50

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

<http://http://docs.sun.com/app/docs/prod/sun.cluster32>

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
820-0335	<i>Sun Cluster 3.2 Documentation Center</i>

TABLE 1 Sun Cluster 3.2 Software Collection for Solaris OS Software Manuals (Continued)

Part Number	Book Title
819-2969	<i>Sun Cluster Concepts Guide for Solaris OS</i>
819-2972	<i>Sun Cluster Data Services Developer's Guide for Solaris OS</i>
819-2974	<i>Sun Cluster Data Services Planning and Administration Guide for Solaris OS</i>
819-2973	<i>Sun Cluster Error Messages Guide for Solaris OS</i>
819-2968	<i>Sun Cluster Overview for Solaris OS</i>
819-6811	<i>Sun Cluster Quick Reference</i>
819-3055	<i>Sun Cluster Reference Manual 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>

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>
819-3058	<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 Service 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 NI Grid Service Provisioning System for Solaris OS</i>
819-0693	<i>Sun Cluster Data Service for NetBackup Guide for Solaris OS</i>
819-2979	<i>Sun Cluster Data Service for NFS 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
819-3061	<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>
819-2983	<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 Service 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>
819-2983	<i>Sun Cluster Data Service for SAP 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>

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

Part Number	Book Title
819-3067	<i>Sun Cluster Data Service for WebSphere MQ Guide for Solaris OS</i>
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.

- “Concepts Guide” on page 50
- “Software Installation Guide” on page 51
- “Sun Cluster Data Services Planning and Administration Guide” on page 51
- “Sun Cluster Data Service for MaxDB Guide” on page 52
- “Sun Cluster Data Service for SAP Guide” on page 52
- “Sun Cluster Data Service for SAP liveCache Guide” on page 53
- “Sun Cluster Data Service for SAP Web Application Server Guide” on page 54
- “System Administration Guide” on page 58
- “Data Services Developer's Guide” on page 61
- “Quorum Server User's Guide” on page 62
- “Man Pages” on page 62

Concepts Guide

This section discusses error and omissions in the *Sun Cluster Concepts Guide for Solaris OS*.

x86: Sun Cluster Topologies for x86

In the section “Sun Cluster Topologies for x86” in *Sun Cluster Concepts Guide for Solaris OS*, the following statement is out of date for the Sun Cluster 3.2 release: “Sun Cluster that is composed of x86 based systems supports two nodes in a cluster.”

The statement should instead read as follows: “A Sun Cluster configuration that is composed of x86 based systems supports up to eight nodes in a cluster that runs Oracle RAC, or supports up to four nodes in a cluster that does not run Oracle RAC.”

Software Installation Guide

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

Missing Upgrade Preparation for Clusters that Run Sun Cluster Geographic Edition Software

If you upgrade a cluster that also runs Sun Cluster Geographic Edition software, there are additional preparation steps you must perform before you begin Sun Cluster software upgrade. These steps include shutting down the Sun Cluster Geographic Edition infrastructure. Go instead to Chapter 4, “Upgrading the Sun Cluster Geographic Edition Software,” in *Sun Cluster Geographic Edition Installation Guide* in *Sun Cluster Geographic Edition Installation Guide*. These procedures document when to return to the Sun Cluster Software Installation Guide to perform Sun Cluster software upgrade.

Sun Cluster Data Services Planning and Administration Guide

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

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.

Sun Cluster Data Service for MaxDB Guide

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

Changes to Sun Cluster Data Service for MaxDB Support on Non-Global Zones on SPARC and x86 Based System

The Sun Cluster Data Service for MaxDB supports non-global zones on SPARC and x86 based systems. The following changes should be made to the Sun Cluster Data Service MaxDB Guide for this support. The following steps can be performed on a cluster that has been configured to run in global zones. If you are installing your cluster to run in non-global zones, a few of these steps might not be necessary as indicated below.

- On each zone, ensure that all of the network resources are present in the `/etc/hosts` file to avoid any failures because of name service lookup.
- On each zone, create an entry for the MaxDB group in the `/etc/group` file, and add potential users to the group.
- On each zone, create an entry for the MaxDB user ID.
Use the following command to update the `/etc/passwd` and `/etc/shadow` files with an entry for the user ID.


```
# useradd -u uid -g group -d /sap-home maxdb user
```
- Create mount point directories in the zones where MaxDB could potentially run.
- Configure the `/etc/nsswitch.conf` file so that Sun Cluster HA for MaxDB starts and stops correctly in the event of a switchover or a failover.
- On each zone update `/etc/services` file with all necessary MaxDB ports obtained from the global zones `/etc/services`. This step might not be necessary for Max DB that is installed in non-global zones.
- Copy `/etc/opt/sdb` from the global zone to all local zone nodes. This step might not be necessary for MaxDB that is being installed in non-global zones.
- Copy `/var/spool/sql` from the global zone to all local zone nodes. This step might not be necessary for MaxDB that is being installed in non-global zones.
- On x86 based systems only, execute `crle -64 -u -l /sapmnt/MaxDBSystemName/exe` on all local zones that will run MaxDB.

Sun Cluster Data Service for SAP Guide

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

Changes to SAP Support on Non-Global Zones on SPARC and x86 Based System

The Sun Cluster Data Service for SAP supports non-global zones on SPARC and x86 based systems. The following changes should be made to the Sun Cluster Data Service SAP Guide for this support. The following steps can be performed on a cluster that has been configured to run in global zones. If you are installing your cluster to run in non-global zones, a few of these steps might not be necessary as indicated below.

- On each zone, ensure that all of the network resources are present in the `/etc/hosts` file to avoid any failures because of name service lookup.
- On each zone, create an entry for the SAP group in the `/etc/group` file, and add potential users to the group.
- On each zone, create an entry for the SAP user ID.

Use the following command to update the `/etc/passwd` and `/etc/shadow` files with an entry for the user ID.

```
# useradd -u uid -g group -d /sap-home sap user
```

- Create mount point directories in the zones where SAP could potentially run.
- Configure the `/etc/nsswitch.conf` file so that Sun Cluster HA for SAP starts and stops correctly in the event of a switchover or a failover.
- On each zone update `/etc/services` file with all necessary SAP ports obtained from the global zones `/etc/services`. This step might not be necessary for SAP that is being installed in non-global zones.
- On x86 based systems only, execute `crle -64 -u -l /sapmnt/SAPSystemName/exe` on all local zones that will run SAP.

Sun Cluster Data Service for SAP liveCache Guide

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

Changes to SAP liveCache Support on Non-Global Zones on SPARC and x86 Based System

The Sun Cluster Data Service for SAP liveCache supports non-global zones on SPARC and x86 based systems. The following changes should be made to the Sun Cluster Data Service SAP liveCache Guide for this support. The following steps can be performed on a cluster that has been configured to run in global zones. If you are installing your cluster to run in non-global zones, a few of these steps might not be necessary as indicated below.

- On each zone, ensure that all of the network resources are present in the `/etc/hosts` file to avoid any failures because of name service lookup.

- On each zone, create an entry for the SAP liveCache group in the `/etc/group` file, and add potential users to the group.
- On each zone, create an entry for the SAP liveCache user ID.
Use the following command to update the `/etc/passwd` and `/etc/shadow` files with an entry for the user ID.

```
# useradd -u uid -g group -d /sap-home sap user
```
- Create mount point directories in the zones where SAP liveCache could potentially run.
- Configure the `/etc/nsswitch.conf` file so that Sun Cluster HA for SAP liveCache starts and stops correctly in the event of a switchover or a failover.
- On each zone update `/etc/services` file with all necessary SAP liveCache ports obtained from the global zones `/etc/services`. This step might not be necessary for SAP liveCache that is being installed in non-global zones.
- Copy `/etc/opt/sdb` from the global zone to all local zone nodes. This step might not be necessary for SAP liveCache that is being installed in non-global zones.
- Copy `/var/spool/sql` from the global zone to all local zone nodes. This step might not be necessary for SAP liveCache that is being installed in non-global zones.
- On x86 based systems only, execute `crle -64 -u -l /sapmnt/SAPSystemName/exe` on all local zones that will run SAP liveCache.

Sun Cluster Data Service for SAP Web Application Server Guide

This section discusses error and omissions in the *Sun Cluster Data Service for SAP Web Application Server Guide for Solaris OS*.

Support for SAP 7.0 for Sun Cluster HA for SAP Web Application Server (6461002)

In SAP 7.0 and NW2004SR1, when a SAP instance is started, the `sapstartsrv` process is started by default. The `sapstartsrv` process is not under the control of Sun Cluster HA for SAP Web Application Server. So, when a SAP instance is stopped or failed over by Sun Cluster HA for SAP Web Application Server, the `sapstartsrv` process is not stopped.

To avoid starting the `sapstartsrv` process when a SAP instance is started by Sun Cluster HA for SAP Web Application, you must modify the `startsap` script. In addition, rename the `/etc/rc3.d/S90sapinit` file to `/etc/rc3.d/xxS90sapinit` on all the Sun Cluster nodes.

Changes to SAP Web Application Server Support on Non-Global Zones on SPARC and x86 Based System

The Sun Cluster Data Service for SAP Web Application Server supports non-global zones on SPARC and x86 based systems. The following changes should be made to the Sun Cluster Data Service SAP Web Application Server Guide for this support. The following steps can be performed on a cluster that has been configured to run in global zones. If you are installing your cluster to run in non-global zones, a few of these steps might not be necessary as indicated below.

- On each zone, ensure that all of the network resources are present in the `/etc/hosts` file to avoid any failures because of name service lookup.
- On each zone, create an entry for the SAP Web Application Server group in the `/etc/group` file, and add potential users to the group.
- On each zone, create an entry for the SAP Web Application Server user ID.
Use the following command to update the `/etc/passwd` and `/etc/shadow` files with an entry for the user ID.


```
# useradd -u uid -g group -d /sap-home sap user
```
- Create mount point directories in the zones where SAP Web Application Server could potentially run.
- Configure the `/etc/nsswitch.conf` file so that Sun Cluster HA for SAP starts and stops correctly in the event of a switchover or a failover.
- On each zone update `/etc/services` file with all necessary SAP ports obtained from the global zones `/etc/services`. This step might not be necessary for SAP Web Application Server that is being installed in non-global zones.
- On x86 based systems only, execute `crle -64 -u -l /sapmnt/SAPSystemName/exe` on all local zones that will run SAP.

Setting Up the SAP Web Application Server on Non-Global Zones for HASP Configuration (6530281)

Use the following procedure to configure a `HASStoragePlus` resource for non-global zones.

Note –

- The entries in the `/etc/vfstab` file for cluster file systems should contain the global keyword in the mount options.
- The SAP binaries that will be made highly available using the `HASStoragePlus` resource should be accessible from the non-global zones.
- In non-global zones, file systems that are used by different resources in different resource groups must reside in a single `HASStoragePlus` resource that resides in a scalable resource group. The `nodelist` of the scalable `HASStoragePlus` resource group must be a superset of the `nodelists` of the application resource groups that have resources which depend on the file systems. These application resources that depend on the file systems must have a strong resource dependency set to the `HASStoragePlus` resource. In addition, the dependent application resource group must have a strong positive resource group affinity set to the scalable `HASStoragePlus` resource group.

▼ How to Set Up the SAP Web Application Server on Non-Global Zones for HASStoragePlus Configuration

- 1 On any node in the cluster, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.
- 2 Create the scalable resource group with non-global zones that contain the `HASStoragePlus` resource.

```
# clresourcegroup create \  
  -p Maximum primaries=m \  
  -p Desired primaries=n \  
  [-n node-zone-list] hasp-resource-group
```

`-p Maximum primaries=m` Specifies the maximum number of active primaries for the resource group.

`-p Desired primaries=n` Specifies the number of active primaries on which the resource group should attempt to start.

`-n node-zone-list` In the node list of a `HASStoragePlus` resource group, specifies the list of `nodename:zonename` pairs as the node list of the `HASStoragePlus` resource group, where the SAP instances can come online.

`hasp-resource-group` Specifies the name of the scalable resource group to be added. This name must begin with an ASCII character.

- 3 Register the resource type for the `HASStoragePlus` resource.

```
# clresourcetype register HASStoragePlus
```

4 Create the HASStoragePlus resource `hasp-resource` and define the SAP filesystem mount points and global device paths.

```
# clresource create -g hasp-resource-group -t SUNW.HASStoragePlus \
  -p GlobalDevicePaths=/dev/global/dsk/d5s2,dsk/d6 \
  -p affinityon=false -p
FilesystemMountPoints=/sapmnt/JSC,/usr/sap/trans,/usr/sap/JSC hasp-resource
```

<code>-g <i>hasp-resource-group</i></code>	Specifies the resource group name.
<code>GlobalDevicePaths</code>	Contains the following values: <ul style="list-style-type: none"> ▪ Global device group names, such as <code>sap-dg</code>, <code>dsk/d5</code> ▪ Paths to global devices, such as <code>/dev/global/dsk/d5s2</code>, <code>/dev/md/sap-dg/dsk/d6</code>
<code>FilesystemMountPoints</code>	Contains the following values: <ul style="list-style-type: none"> ▪ Mount points of local or cluster file systems, such as <code>/local/mirrlogA</code>, <code>/local/mirrlogB</code>, <code>/sapmnt/JSC</code>, <code>/usr/sap/JSC</code>

The HASStoragePlus resource is created in the enabled state.

5 Register the resource type for the SAP application.

```
# clresourcetype register resource-type
resource-type Specifies the name of the resource type to be added. For more information, see
“Supported Products” on page 21.
```

6 Create a SAP resource group.

```
# clresourcegroup create [-n node-zone-list] -p
RG_affinities=++hastorageplus-rg resource-group-1
resource-group-1 Specifies the SAP services resource group.
```

7 Add the SAP application resource to `resource-group-1` and set the dependency to `hastorageplus-1`.

```
# clresource create -g resource-group-1 -t SUNW.application \
  [-p "extension-property[{node-specifier}]"=value, ?] \
  -p Resource_dependencies=hastorageplus-1 resource
```

8 Bring the failover resource group online.

```
# clresourcegroup online resource-group-1
```

System Administration Guide

This section discusses error and omissions in the *Sun Cluster System Administration Guide for Solaris OS*.

Taking a Solaris Volume Manager Metaset From Nodes Booted in Non-Cluster Mode

▼ How to Take a Solaris Volume Manager Metaset From Nodes Booted in Non-Cluster Mode

Use this procedure to run an application outside the cluster for testing purposes.

- 1 Determine if the quorum device is used in the Solaris Volume Manager metaset, and determine if the quorum device uses scsi2 or scsi3 reservations.

```
# clquorum show
```

- a. If the quorum device is in the Solaris Volume Manager metaset, add a new quorum device which is not part of the metaset to be taken later in non-cluster mode.

```
# clquorum add did
```

- b. Remove the old quorum device.

```
# clquorum remove did
```

- c. If the quorum device uses a scsi2 reservation, scrub the scsi2 reservation from the old quorum and verify that there are no scsi2 reservations remaining.

```
# /usr/cluster/lib/sc/pgre -c pgre_scrub -d /dev/did/rdisk/dids2
```

```
# /usr/cluster/lib/sc/pgre -c pgre_inkeys -d /dev/did/rdisk/dids2
```

- 2 Evacuate the node you want to boot in non-cluster mode.

```
# clresourcegroup evacuate -n targetnode
```

- 3 Take offline any resource group or resource groups that contain HAStorage or HAStoragePlus resources and contain devices or file systems affected by the metaset you want to later take in non-cluster mode.

```
# clresourcegroup offline resourcegroupname
```

- 4 Disable all the resources in the resource groups you took offline.

```
# clresource disable resourcename
```

- 5 **Unmanage the resource groups.**
`# clresourcegroup unmanage resourcegroupname`
- 6 **Take offline the corresponding device group or device groups.**
`# cldevicegroup offline devicegroupname`
- 7 **Disable the device group or device groups.**
`# cldevicegroup disable devicegroupname`
- 8 **Boot the passive node into non-cluster mode.**
`# reboot -x`
- 9 **Verify that the boot process has completed on the passive node before proceeding.**
 - Solaris 9
The login prompt will only appear after the boot process has completed, so no action is required.
 - Solaris 10

`# svcs -x`
- 10 **Determine if there are any scsi3 reservations on the disks in the metaset or metasets. Perform the following commands on all disks in the metasets.**
`# /usr/cluster/lib/sc/scsi -c inkeys -d /dev/did/rdisk/dids2`
- 11 **If there are any scsi3 reservations on the disks, scrub them.**
`# /usr/cluster/lib/sc/scsi -c scrub -d /dev/did/rdisk/dids2`
- 12 **Take the metaset on the evacuated node.**
`# metaset -s name -C take -f`
- 13 **Mount the filesystem or filesystems containing the defined device on the metaset.**
`# mount device mountpoint`
- 14 **Start the application and perform the desired test. After finishing the test, stop the application.**
- 15 **Reboot the node and wait until the boot process has finished.**
`# reboot`
- 16 **Bring online the device group or device groups.**
`# cldevicegroup online -e devicegroupname`

17 Start the resource group or resource groups.

```
# clresourcegroup onLine -eM resourcegroupname
```

Using Solaris IP Filtering with Sun Cluster

Sun Cluster supports Solaris IP Filtering with the following restrictions:

- Only failover data services are supported.
Sun Cluster does not support IP Filtering with scalable data services.
- Only stateless filtering is supported.
- NAT routing is not supported.
- Use of NAT for translation of local addresses is supported. NAT translation rewrites packets on-the-wire and is therefore transparent to the cluster software.

▼ How to Set Up Solaris IP Filtering**1 In the `/etc/iu.ap` file, modify the public NIC entries to list `clhbsndrpfil` as the module list.**

The `pfil` must be the last module in the list.

Note – If you have the same type of adapter for private and public network, your edits to the `/etc/iu.ap` file will push `pfil` to the private network streams. However, the cluster transport module will automatically remove all unwanted modules at stream creation, so `pfil` will be removed from the private network streams.

2 To ensure that the IP filter works in non-cluster mode, update the `/etc/ipf/pfil.ap` file.

Updates to the `/etc/iu.ap` file are slightly different. See the IP Filter documentation for more information.

3 Reboot all affected nodes.

You can boot the nodes in a rolling fashion.

4 Add filter rules to the `/etc/ipf/ipf.conf` file on all affected nodes. For information on IP filter rules syntax, see `ipf(4)`

Keep in mind the following guidelines and requirements when you add filter rules to Sun Cluster nodes.

- Sun Cluster fails over network addresses from node to node. No special procedure or code is needed at the time of failover.
- All filtering rules that reference IP addresses of logical hostname and shared address resources must be identical on all cluster nodes.

- Rules on a standby node will reference a non-existent IP address. This rule is still part of the IP filter's active rule set and will become effective when the node receives the address after a failover.
 - All filtering rules must be the same for all NICs in the same IPMP group. In other words, if a rule is interface-specific, the same rule must also exist for all other interfaces in the same IPMP group.
- 5 **Enable the `ipfilter` SMF service.**
- ```
svcadm enable /network/ipfilter:default
```

## 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 Reconfiguration 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.

## 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, omissions, and additions in the Sun Cluster man pages.

### ccp(1M)

The following revised Synopsis and added Options sections of the ccp(1M) man page document the addition of Secure Shell support to the Cluster Control Panel (CCP) utilities:

## SYNOPSIS

```
$CLUSTER_HOME/bin/ccp [-s] [-l username] [-p ssh-port] {clustername | nodename}
```

## OPTIONS

The following options are supported:

- l *username*** Specifies the user name for the ssh connection. This option is passed to the `cconsole`, `crlogin`, or `cssh` utility when the utility is launched from the CCP. The `ctelnet` utility ignores this option.

If the `-l` option is not specified, the user name that launched the CCP is effective.
- p *ssh-port*** Specifies the Secure Shell port number to use. This option is passed to the `cssh` utility when the utility is launched from the CCP. The `cconsole`, `crlogin`, and `ctelnet` utilities ignore this option.

If the `-p` option is not specified, the default port number 22 is used for secure connections.
- s** Specifies using Secure Shell connections to node consoles instead of `telnet` connections. This option is passed to the `cconsole` utility when the utility is launched from the CCP. The `crlogin`, `cssh`, and `ctelnet` utilities ignore this option.

If the `-s` option is not specified, the `cconsole` utility uses `telnet` connections to the consoles.

To override the `-s` option, deselect the Use SSH checkbox in the Options menu of the `cconsole` graphical user interface (GUI).

## **cconsole(1M), crlogin(1M), cssh(1M), and ctelnet(1M)**

The following revised Synopsis and added Options sections of the combined `cconsole`, `crlogin`, `cssh`, and `ctelnet` man page document the addition of Secure Shell support to the Cluster Control Panel utilities:

## SYNOPSIS

```
$CLUSTER_HOME/bin/cconsole [-s] [-l username] [clustername... | nodename...]
$CLUSTER_HOME/bin/crlogin [-l username] [clustername... | nodename...]
$CLUSTER_HOME/bin/cssh [-l username] [-p ssh-port] [clustername... | nodename...]
$CLUSTER_HOME/bin/ctelnet [clustername... | nodename...]
```

## DESCRIPTION

`cssh` This utility establishes Secure Shell connections directly to the cluster nodes.

### OPTIONS

- `-l username` Specifies the `ssh` user name for the remote connections. This option is valid with the `cconsole`, `crlogin`, and `cssh` commands.
- The argument value is remembered so that clusters and nodes that are specified later use the same user name when making connections.
- If the `-l` option is not specified, the user name that launched the command is effective.
- `-p ssh-port` Specifies the Secure Shell port number to use. This option is valid with the `cssh` command.
- If the `-p` option is not specified, the default port number 22 is used for secure connections.
- `-s` Specifies using Secure Shell connections instead of `telnet` connections to node consoles. This option is valid with the `cconsole` command.
- If the `-s` option is not specified, the utility uses `telnet` connections to the consoles.
- To override the `-s` option from the `cconsole` graphical user interface (GUI), deselect the Use SSH checkbox in the Options menu.

### `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 operates 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.

### **serialports(4)**

The following information is an addition to the `Description` section of the `serialport(4)` man page:

To support Secure Shell connections to node consoles, specify in the `/etc/serialports` file the name of the console-access device and the Secure Shell port number for each node. If you use the default Secure Shell configuration on the console-access device, specify port number 22.

### **SUNW.Event(5)**

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

