

Sun Cluster 3.2 11/09 Release Notes

3.2 11/09 Release Notes

821-0258

Release Notes

Copyright © 2007, 2009, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related software documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications which may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark licensed through X/Open Company, Ltd.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

(English) Sun Cluster 3.2 11-09 Release Notes

This document provides the following information for Sun™ Cluster 3.2 11/09 software.

- [What's New in Sun Cluster 3.2 11/09 Software](#)
- [Features Nearing End of Life](#)
- [Compatibility Issues](#)
- [Commands Modified in This Release](#)
- [Product Name Changes](#)
- [Supported Products](#)
- [Sun Cluster Security Hardening](#)
- [Known Issues and Bugs](#)
- [Patches and Required Firmware Levels](#)
- [Sun Cluster 3.2 11/09 Documentation](#)
- [Localization](#)
- [Documentation Issues](#)

[Top](#)

What's New in Sun Cluster 3.2 11/09 Software

This section provides information related to new features, functionality, and supported products in Sun Cluster 3.2 11/09 software.

The following new features and functionality are provided in patches to Sun Cluster 3.2 11/09 software:

- [CLI Support to Configure Oracle RAC Instance Proxy Resource Group](#)
- [HA for Oracle Support for Third-Party Volume Managers with Clustered Oracle ASM Disk Groups](#)

The following new features and functionality are provided in the initial Sun Cluster 3.2 11/09 release:

- [Solaris Volume Manager Three-Mediator Support](#)
- [Oracle Automatic Storage Manager \(ASM\) on HA-Oracle and Oracle RAC](#)
- [SWIFTAlliance Access and Gateway 6.3](#)
- [Outgoing Connection Support For Scalable Services](#)
- [IPsec Security Association \(SA\) Failover for Scalable Services](#)
- [Campus Cluster With Hitachi Universal Replicator](#)
- [Sun Cluster HA for LDoms Guest Domains](#)
- [Generic Affinity for Scalable Services](#)
- [Round-Robin Load-Balancing Scheme for Scalable Services](#)
- [SCTP Support in Scalable Services](#)
- [Maintaining Solaris Systems with Solaris Live Upgrade and Update on Attach](#)
- [Sun Storage J4200/J4400 SATA Arrays](#)

[Top](#)

CLI Support to Configure Oracle RAC Instance Proxy Resource Group

Support is added to the Sun Cluster maintenance commands to configure an Oracle Real Application Clusters (Oracle RAC) instance proxy resource group when Oracle ASM is not also configured. This support is delivered starting in the following HA-Oracle patch versions:

- SPARC: 126047-16
- x86: 126048-16

For supporting configuration instructions, see [Configuring an Oracle RAC Instance Proxy Resource Group Without ASM \(CLI\)](#) .

Procedures for removing the Grid Infrastructure resource are also provided. See [Removing an Oracle Grid Infrastructure Resource](#).

[Top](#)

HA for Oracle Support for Third-Party Volume Managers with Clustered Oracle ASM Disk Groups

Support is added to the HA for Oracle data service for third-party volume managers configured with clustered Oracle ASM disk groups. This support is delivered in the following HA for Oracle patches:

- SPARC: 126047-16
- x86: 126048-16

For supporting configuration instructions, see [Support for Third-Party Volume Managers with Clustered Oracle ASM Disk Groups](#).

[Top](#)

Solaris Volume Manager Three-Mediator Support

Sun Cluster software now supports mediators for three-room campus cluster configurations that use Solaris Volume Manager or multi-owner Solaris Volume Manager for Sun Cluster. A two-room (two-node) campus cluster now works with a third mediator host outside the cluster. The third mediator host does not have to be attached to the shared storage that contains the disk set for which the host is a mediator.

The mediator host uses Solaris Volume Manager to facilitate automatic recovery for a two-room campus cluster by tracking which mirrored half of the storage is the most up to date. The third mediator then provides mediator quorum to allow Solaris Volume Manager to recover from a destroyed room.

See [Changes to Guidelines for Mediators](#) for the revised guidelines for configuring mediators.

[Top](#)

Oracle Automatic Storage Manager (ASM) on HA-Oracle and Oracle RAC

This release provides support for HA-Oracle with Oracle's Automated Storage Management (ASM) product, and also extends Sun Cluster support for Oracle RAC to allow management of ASM within the Sun Cluster RAC management framework. For more information, see the [Sun Cluster Data Service for Oracle Guide](#) or the [Sun Cluster Data Service for Oracle RAC Guide](#).



Note

Additional information is provided [Oracle ASM for Sun Cluster Support for RAC](#).

[Top](#)

SWIFTAlliance Access and Gateway 6.3

Sun Cluster 3.2 11/09 software supports SWIFTAlliance Access 6.3 and SWIFTAlliance Gateway 6.3 from SWIFT with the corresponding data services.

[Top](#)

Outgoing Connection Support for Scalable Services

Traditional scalable services only handle requests that originate from clients outside the cluster. Outgoing requests that originate from the scalable service carry the IP address of the cluster node but not the virtual IP address that is provided by the `SUNW.SharedAddress` resource. The outgoing connection support feature enables a scalable service to initiate requests to a server that is outside the cluster, and to use the virtual IP address as the source IP address of those requests. This hides the IP addresses of the individual cluster nodes and allows for better abstraction of network topology and more flexible setup of network features, such as filtering and security between clusters that belong to different organizations.

For more information, see [Outgoing Connection Support for Sun Cluster Scalable Services](#).

[Top](#)

IPsec Security Association (SA) Failover for Scalable Services

In a scalable service setup using IPsec, clients must re-negotiate the IPsec Security Association (SA) state after a failover, since all SA session states are specific to a node. This increases the recovery time that is perceived by clients and produces an unnecessary processing load on the cluster. The IPsec Security Association feature enables some amount of replication of SA state between cluster nodes, to eliminate the need for SA re-negotiation after a failover.

For more information, see [IPsec Security Association Failover for Sun Cluster Scalable Services](#).

[Top](#)

Campus Cluster With Hitachi Universal Replicator

Support in Sun Cluster 3.2 11/09 software for Hitachi TrueCopy has been extended to Hitachi Universal Replicator. When Hitachi Universal Replicator is deployed in a campus cluster, setup and management within the cluster is the same as it is for Hitachi TrueCopy. For more information, see the [Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy and Universal Replicator](#).

[Top](#)

Sun Cluster HA for LDOMs Guest Domain

Sun Cluster software now supports a highly available data service that provides fault monitoring and failover for an LDOMs guest domain. For more information, see [Sun Cluster HA for LDOMs Guest Domain](#).

[Top](#)

Generic Affinity for Scalable Services

The `LB_STICKY` or `LB_STICKY_WILD` weighting policies set IP affinity for scalable services. Current Sun Cluster implementation provides IP affinity only for TCP services. No IP affinity is guaranteed if a service group has only UDP service. The generic affinity feature provides IP affinity to the client, irrespective of the protocol. A new data service property, `generic_affinity` provides this functionality.

Details are available at [Generic Affinity for Sun Cluster Scalable Services](#).

[Top](#)

Round-Robin Load-Balancing Scheme for Scalable Services

Current Sun Cluster scalable services provide random load balancing. This new round-robin load-balancing scheme provides a balanced distribution of load, regardless of the number of requests and clients. It introduces three new resource properties for the scalable data service configuration: `Round_Robin`, which enables the load-balancing scheme, `Conn_threshold`, which cleans up inactive TCP connections, and `Udp_session_timeout` which cleans up inactive UDP connections.

Details are available at [Round-Robin Load-Balancing Scheme for Sun Cluster Scalable Services](#).

[Top](#)

SCTP Support in Scalable Services

Sun Cluster 3.2 11/09 software introduces support for a new transport protocol called Stream Control Transmission Protocol (SCTP). The `port_list` resource property can now include the value `sctp`, in addition to `udp` and `tcp`.

For more information, see [SCTP Support in Sun Cluster Scalable Services](#).

[Top](#)

Maintaining Solaris Systems with Solaris Live Upgrade and Update on Attach

Solaris Live Upgrade and the Solaris Update on attach feature can be used with the Sun Cluster 3.2 software to minimize the planned downtime associated with software installation and upgrades. Descriptions of these functions and detailed directions for their use with the Sun Cluster software are available at [Maintaining Solaris with Live Upgrade and Update on Attach](#).



Note:

Using update on attach to install Sun Cluster patches is currently not supported.

Previous releases of the Sun Cluster software did not support using Live Upgrade with Sun Cluster systems that boot from a ZFS root file system. A fix to this limitation has been integrated in the Sun Cluster 3.2 11/09 release.

[Top](#)

Support for Sun Storage J4200/J4400 SATA Arrays

Sun Cluster 3.2 11/09 software now supports the use of Sun Storage J4200 and J4400 arrays with SATA HDDs in a cluster configuration. Sun Storage J4200 and J4400 storage arrays with SATA drives rely on an internal STP/SATA bridge, which handles the mapping to SATA. Storage arrays with SATA drives have certain limitations that must be taken into consideration in a Sun Cluster configuration. Observe the following requirements and guidelines when you configure dual-ported Sun Storage J4200 and J4400 SATA drives for use in a Sun Cluster configuration:

- **SATA constraint** - Because of the SATA protocol's point-to-point nature, only one initiator can be affiliated with a SATA drive on each J4200 or J4400 array's SIM card.
- **Cabling**- To work around the SATA protocol constraints, the cabling layout from hosts to the J4200 or J4400 arrays might expose the SIMs and the HBAs as a single point of failure in certain configurations. Figure 2 shows an example of such a configuration.
- **Initiators** - You must use one initiator from a host to a SIM card, as shown in Figures 1 and 2.

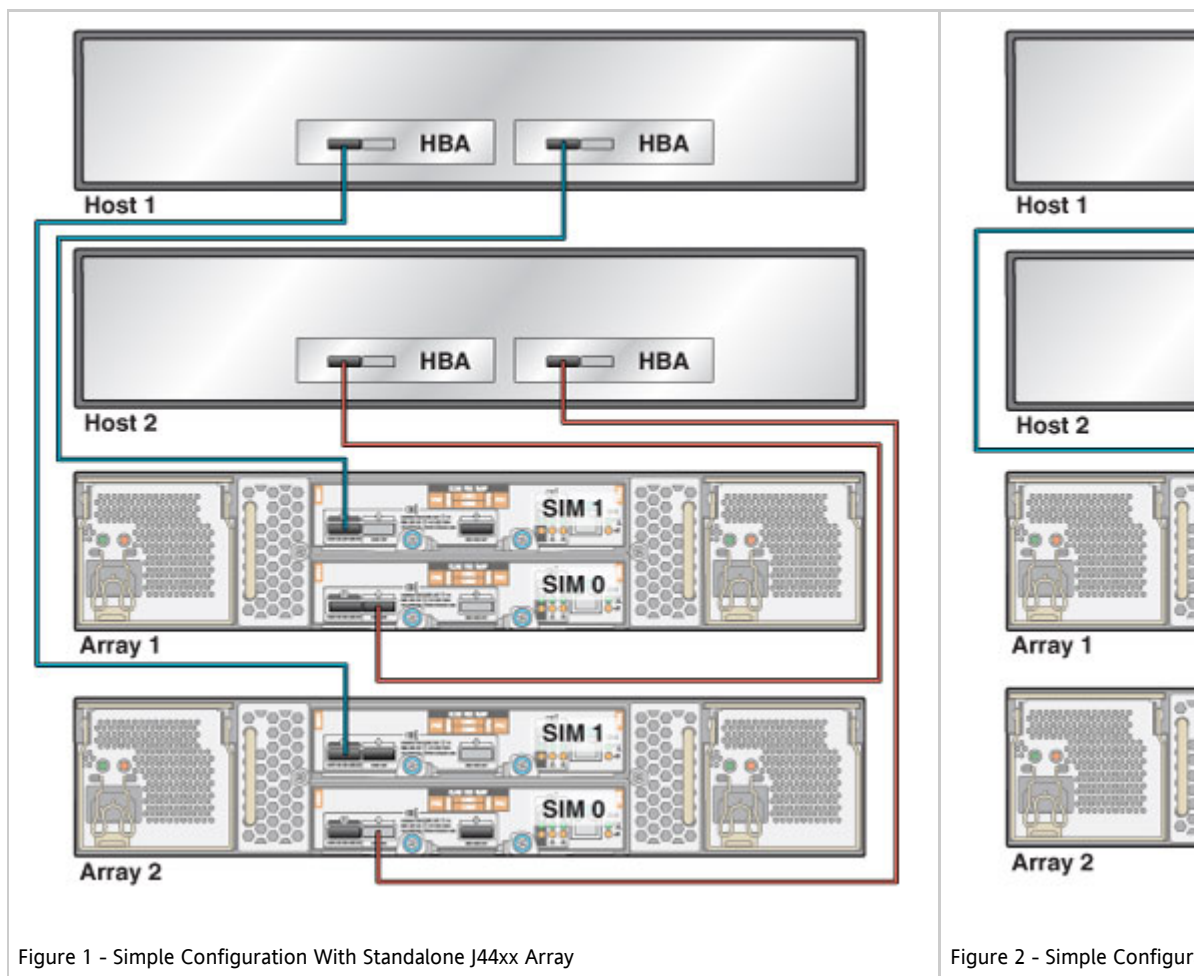


Figure 1 - Simple Configuration With Standalone J44xx Array

Figure 2 - Simple Configuration

- **Mirroring** - To avoid possible loss of data if a disk fails, use host-based mirroring of the data.
- **Sharing arrays** - You cannot share a SATA JBOD array between hosts that belong to different clusters.

- Disabling global fencing- You must disable global fencing for all disks in the SATA storage array. This is required regardless of whether any SATA disks are to be used as a quorum device or simply as shared storage.
 - To disable fencing during initial cluster configuration, run the `scinstall` utility in Custom Mode. When prompted by the utility whether to disable global fencing, respond "Yes".
 - To disable fencing of all storage devices in the cluster after the cluster is established, follow instructions in "How to Change the Default Global Fencing Protocol Settings for All Storage Devices" in the Sun Cluster System Administration Guide to set the `global_fencing` property to `nofencing_noscrub`. Alternatively, follow instructions in "How to Change the Fencing Protocol for a Single Storage Device" to set the `default_fencing` property of each SATA disk to `nofencing_noscrub`.

For more information about setting up Sun Storage J4200 or J4400 arrays, see the Sun Storage J4200/J4400 Array System Overview at <http://docs.sun.com/app/docs/doc/820-3223-14> and the Sun Storage J4200/J4400 Array System Hardware Installation Guide at <http://docs.sun.com/app/docs/doc/820-3218-14>.

[Top](#)

Features Nearing End of Life

The following feature is nearing end of life in Sun Cluster software.

Original Administrative Command Set

The original Sun Cluster administrative command set, including commands with the prefix `sc`, might no longer be supported in a future release. Only the object-oriented command set, including commands beginning with the prefix `cl`, will be available.

[Top](#)

Compatibility Issues

- [Accessibility Features for People With Disabilities](#)
- [Solaris Volume Manager GUI](#)
- [Live Upgrade Fails in Clusters Configured With Zone Clusters](#)
- [Sun LDoms Warm Migration Is Not Supported With the HA-LDoms Data Service](#)
- [Oracle Database 11.2.0.2](#)

This section contains information about Sun Cluster compatibility issues.

- Additional Sun Cluster framework compatibility issues are documented in "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 Upgrade Guide for Solaris OS.
- For other known problems or restrictions, see [Known Issues and Bugs](#).

[Top](#)

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 on request to determine which versions are best suited for deploying accessible solutions.

[Top](#)

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.

[Top](#)

Live Upgrade Fails in Clusters Configured With Zone Clusters (6852390)

Problem Summary: On a cluster that is configured with a zone cluster, the use of Live Upgrade to upgrade the global cluster fails. The Live Upgrade scripts use an underlying `zoneadm move` command to move zones to the correct paths when switching to the new boot environment (BE). But this command is unable to perform this move for the `cluster` branded non-global zones of a zone cluster.

Workaround: To upgrade a cluster that has one or more zone clusters configured in it, use an upgrade method other than Live Upgrade. Contact your Sun service representative to determine whether a patch might be available.

[Top](#)

Sun LDom's Warm Migration Is Not Supported With the HA-LDom's Data Service (6882083)

Problem Summary: When warm migration is performed on a Solaris Logical Domains (LDoms) guest domain that is configured with the HA-LDom's data service, the guest domain OS does not boot after the Sun Cluster switchover.

Workaround: This issue is resolved in LDoms version 1.3. Do not configure the HA-LDom's data service with LDoms version 1.2 or earlier if you want to use the LDoms warm migration feature.

[Top](#)

Oracle Database 11.2.0.2

For Oracle Database version 11.2.0.2, when using the Oracle Universal Installer (OUI) to install Oracle Grid Infrastructure, you must specify the Oracle ASM discovery string so that only devices that are accessible by the Oracle Grid Infrastructure user will match. If other devices are matched that the Oracle Grid Infrastructure user cannot access, the installation fails.

For example, instead of setting the disk discovery path as `/dev/did/rdisk/d*`, list only the specific disks that are accessible by the Oracle Grid Infrastructure software owner, such as `/dev/did/rdisk/d1s5`, `/dev/did/rdisk/d2s5`.

For more information, see <https://support.us.oracle.com/oip/faces/secure/km/DocumentDisplay.jspx?id=1303209.1>.

[Top](#)

Commands Modified in This Release

There are no changes to the Sun Cluster command interfaces in this release that might cause user scripts to fail.

[Top](#)

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 11/09 software is distributed under Solaris Cluster 3.2 11/09 and Sun Java Availability Suite.

[Top](#)

Current Product Name	Former Product Name
Sun Cluster HA for MaxDB	Sun Cluster HA for SAP DB

Sun QFS	Sun StorEdge QFS, Sun StorageTek QFS
Sun Ops Center	Sun xVM Ops Center

[Top](#)

Supported Products

This section describes the supported software and memory requirements for Sun Cluster 3.2 11/09 software.

- [Data Services](#)
- [File Systems](#)
- [Memory Requirements](#)
- [Solaris Operating System \(OS\)](#)
- [Sun Logical Domains \(LDMs\)](#)
- [Sun Management Center](#)
- [Sun StorageTek Availability Suite](#)
- [Sun StorEdge Availability Suite](#)
- [Volume Managers](#)

[Top](#)

Data Services

Contact your Sun sales representative for the complete list of supported data services (agents) and application versions. Links to the documentation for many data service agents are available [here](#).



Note

The Sun Cluster Data Service for Oracle Guide is translated into Japanese and Simplified Chinese. Documentation for other data services, including man pages and wizard online help, is no longer translated from English to other languages.

[Top](#)

File Systems

Solaris 10 SPARC

File System	Additional Information
Solaris UFS	
Solaris ZFS	Not supported for the <code>/globaldevices</code> file system
Sun StorEdge QFS	
Sun QFS 5.0 - Standalone file system	Supported Data Services: All failover data services External Volume Management: Solaris Volume Manager only
QFS 4.6 - Standalone file system	Supported Data Services: All failover data services External Volume Management: Solaris Volume Manager, VxVM
QFS 4.6 and 5.0 - Shared QFS file system	Supported Data Services: Oracle RAC External Volume Management: Solaris Volume Manager for Sun Cluster

QFS 4.6 and 5.0 - Shared QFS clients outside the cluster (SC-COTC)	Supported Data Services: None; only a shared file system is supported External Volume Management: No external volume manager is supported
QFS 4.6 and 5.0 - HA-SAM Failover	Supported Data Services None; only a shared file system is supported External Volume Management: No external volume manager is supported
Veritas File System components that are delivered as part of Veritas Storage Foundation 5.0 and 5.1. Support for VxFS 5.0 requires VxFS 5.0 MP3.	

[Top](#)

Solaris 10 x86

File System	Additional Information
Solaris UFS	
Solaris ZFS	Not supported for the <code>/globaldevices</code> file system
Sun StorEdge QFS	
Sun QFS 5.0 - Standalone file system	Supported Data Services: All failover data services External Volume Management: Solaris Volume Manager only
QFS 4.6 - Standalone file system	Supported Data Services: All failover data services External Volume Management: Solaris Volume Manager, VxVM
QFS 4.6 and 5.0 - Shared QFS file system	Supported Data Services: Oracle RAC External Volume Management: Solaris Volume Manager for Sun Cluster
QFS 4.6 and 5.0 - Shared QFS clients outside the cluster (SC-COTC)	Supported Data Services: None; only a shared file system is supported External Volume Management: No external volume manager is supported
QFS 4.6 and 5.0 - HA-SAM Failover	Supported Data Services None; only a shared file system is supported External Volume Management: No external volume manager is supported
Veritas File System components that are delivered as part of Veritas Storage Foundation 5.0 and 5.1. Support for VxFS 5.0 requires VxFS 5.0 MP3.	

[Top](#)

Solaris 9 SPARC

File System	Additional Information
Solaris UFS	
Sun StorEdge QFS	
Sun QFS 5.0 - Standalone file system	Supported Data Services: All failover data services External Volume Management: Solaris Volume Manager only

QFS 4.6 - Standalone file system	Supported Data Services: All failover data services External Volume Management: Solaris Volume Manager, VxVM
QFS 4.6 and 5.0 - Shared QFS file system	Supported Data Services: Oracle RAC External Volume Management: Solaris Volume Manager for Sun Cluster
QFS 4.6 and 5.0 - Shared QFS clients outside the cluster (SC-COTC)	Supported Data Services: None; only a shared file system is supported External Volume Management: No external volume manager is supported
QFS 4.6 and 5.0 - HA-SAM Failover	Supported Data Services None; only a shared file system is supported External Volume Management: No external volume manager is supported
Veritas File System components that are delivered as part of Veritas Storage Foundation 5.0 and 5.1. Support for VxFS 5.0 requires VxFS 5.0 MP3.	

[Top](#)

Memory Requirements

Sun Cluster 3.2 11/09 software requires the following memory requirements for every cluster node:

- Minimum of 1 Gbytes of physical RAM (2 Gbytes typical)
- Minimum of 6 Gbytes 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.

[Top](#)

Solaris Operating System (OS)

Sun Cluster 3.2 11/09 software and Quorum Server software require the following minimum versions of the Solaris OS:

- Solaris 9 (SPARC only) - Solaris 9 9/05, Solaris 9 9/05 HW
- Solaris 10 - Solaris 10 5/09, Solaris 10 10/09



Note

Sun Cluster software does not support multiple versions of Solaris software in the same running cluster.

[Top](#)

Sun Logical Domains (LDMs)

This Sun Cluster release supports Sun Logical Domains 1.2 software on the Solaris 10 OS only. A minimum of patch 142840-02 is required.

[Top](#)

Sun Management Center

This Sun Cluster release supports Sun Management Center software versions 3.6.1 and 4.0.

[Top](#)

Sun StorageTek™ Availability Suite

This Sun Cluster release supports Sun StorageTek Availability Suite 4.0 software on the Solaris 10 OS only.

[Top](#)

Sun StorEdge™ Availability Suite

This Sun Cluster release supports Sun StorEdge Availability Suite 3.2.1 software on the Solaris 9 OS only.

[Top](#)

Volume Managers

This Sun Cluster release supports the following volume managers.

Solaris 10 SPARC

Volume Manager	Cluster Feature
Solaris Volume Manager	Solaris Volume Manager for Sun Cluster
Veritas Volume Manager components that are delivered as part of Veritas Storage Foundation 5.0 and 5.1. VxVM 5.0 support requires VxVM 5.0 MP3 RP2.	VxVM 5.0 MP3 RP2 and VxVM 5.1 cluster feature (with RAC only)

[Top](#)

Solaris 10 x86

Volume Manager	Cluster Feature
Solaris Volume Manager	Solaris Volume Manager for Sun Cluster
Veritas Volume Manager components that are delivered as part of Veritas Storage Foundation 5.0 and 5.1. VxVM 5.0 support requires VxVM 5.0 MP3 RP2. Veritas Volume Replicator (VVR) and Fast Mirror Resynchronization (FMR) are not supported on the x86 platform.	Not applicable - Sun Cluster 3.2 11/09 software does not support the VxVM cluster feature on the x86 platform.

[Top](#)

Solaris 9 SPARC

Volume Manager	Cluster Feature
Solaris Volume Manager	Solaris Volume Manager for Sun Cluster
Veritas Volume Manager components that are delivered as part of Veritas Storage Foundation 5.0 and 5.1. VxVM 5.0 support requires VxVM 5.0 MP3 RP2.	VxVM 5.0 MP3 RP2 and VxVM 5.1 cluster feature (with RAC only)

[Top](#)

Sun Cluster Security Hardening

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

Cluster Security Hardening.

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

- Sun Cluster HA for Apache
- Sun Cluster HA for Apache Tomcat
- Sun Cluster HA for DHCP
- Sun Cluster HA for DNS
- Sun Cluster HA for MySQL
- Sun Cluster HA for N1 Grid Engine
- Sun Cluster HA for NFS
- Sun Cluster HA for Oracle
- Sun Cluster HA for Oracle E-Business Suite
- 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 WebLogic Server
- Sun Cluster HA for WebSphere MQ
- Sun Cluster HA for WebSphere MQ Integrator

[Top](#)

Known Issues and Bugs

The following known issues and bugs affect the operation of the Sun Cluster 3.2 11/09 release. Bugs and issues are grouped into the following categories:

- [Administration](#)
- [Data Services](#)
- [Installation](#)
- [Localization](#)
- [Runtime](#)
- [Upgrade](#)

[Top](#)

Administration

Oracle's SPARC T3-4 Fails During Reboot (6993321)

Problem Summary: During a reboot, Oracle's SPARC T3-4 server with four processors fails to connect to the Oracle Solaris Cluster framework. For more detailed information, see 6986241. Error messages similar to the following appear:

```

Sep 20 15:18:53 svc.startd [8]: svc:/system/pools:default: Method or service exit timed out. Killing contract 29.
Sep 20 15:18:53 svc.startd [8]: svc:/system/pools:default: Method "/lib/svc/method/svc-pools start" failed due to
signal KILL.
...
Sep 20 15:20:55 solta svc.startd [8]: system/pools:default failed: transitioned to maintenance (see 'svcs -xv' for
details)
...
Sep 20 15:22:12 solta INITGCHB: Given up waiting for rgmd.
...
Sep 20 15:23:12 solta Cluster.GCHB_resd: GCHB system error: scha_cluster_open failed with 18
Sep 20 15:23:12 solta : No such process

```

Workaround: Use the `svccfg` command to increase the service timeout to 300 seconds. Boot into non-cluster mode and perform the following commands:

```

svccfg -s svc:/system/pools setprop start/timeout_seconds = 300
svcadm refresh svc:/system/pools

```

After you perform these commands, boot into cluster mode.

[Top](#)

Removing Nodes from the Cluster Configuration Can Result in Node Panics (6735924)

Problem Summary: Changing a cluster configuration from a three-node cluster to a two-node cluster might result in complete loss of the cluster, if one of the remaining nodes leaves the cluster or is removed from the cluster configuration.

Workaround: Immediately after removing a node from a three-node cluster configuration, run the `cldevice clear` command on one of the remaining cluster nodes.

[Top](#)

Global-Cluster Node Cannot Get Zone Cluster Resource Group or Resource Information If Its Zone-Cluster Node Has Never Been Booted Into Cluster Mode (6770391)

Problem Summary: This problem occurs for a zone cluster that is configured with resources and resource groups. If a zone-cluster node has not booted into cluster mode even once, then running the `clresourcegroup status -Z zone_cluster_name` and `clresource status -Z zone_cluster_name` commands in the global zone returns an unexpected error message. Similarly, running the `scha_resourcegroup_get -Z zone_cluster_name` and `scha_resource_get -Z zone_cluster_name` commands in the global zone returns error 18, if the zone-cluster nodes are not yet booted into cluster mode. On such nodes, Sun Cluster Manager does not report the status of resource groups and resources for the zone cluster.

Workaround: Boot the zone-cluster nodes into cluster mode by running the following command in the global zone, where `nodelist` is the comma-separated list of nodes on which the zone-cluster nodes are not yet booted into cluster mode:

```
# clzonecluster boot -n nodelist zone_cluster_name
```

This command starts the cluster services that are required to work around this problem. You only need to boot the zone-cluster nodes once when the physical node is booted. After that, you can halt or boot the zone-cluster nodes at your discretion, for as long as the physical node remains up.

[Top](#)

First Attempt to Remove the RAC Framework Resource Group From the Cluster Using `clsetup` Fails (6876092)

Problem Summary: The menu-driven `clsetup` command provides an option to delete an existing RAC framework resource

group from the cluster. The task fails the first time it is invoked. The following error message is displayed when the task is attempted:

```
# /usr/cluster/bin/clresource disable -R -g rac-framework-rg +
# /usr/cluster/bin/clresource delete-g rac-framework-rg +
# /usr/cluster/bin/clresourcegroup delete rac-framework-rg
clresourcegroup: (C667636) rac-framework-rg: resource group is undergoing a
reconfiguration, try again later

Command failed.

Press Enter to continue:
```

Workaround: There are two workarounds for this problem:

Option 1: Use the non-interactive command-line interface to delete the RAC framework resource group. Run the following commands at the command line:

```
# clresource disable -R -g rac-framework-rg +
# clresource delete -g rac-framework-rg +
# clresourcegroup delete rac-framework-rg
```

Option 2: Rerun the `clsetup` option to remove the RAC framework resource group. The second attempt should succeed.

[Top](#)

Verify Hitachi TrueCopy and EMC Symmetrix Remote Data Facility (SRDF) Replica States Before Switching Replicated Device Groups to Another Node (6798901)

Problem Summary: If a Hitachi TrueCopy device group replica pair is in the `COPY` state, or if a Symmetrix Remote Data Facility (SRDF) device group replica pair is split, attempts to switch the device group over to another node fails. Furthermore, the device group is unable to come back online on the original node until the replica pair has been returned to a paired state.

Workaround: Verify that TrueCopy replicas are not in the `COPY` state, or that SRDF replicas are not split, before you attempt to switch the associated Sun Cluster device group to another cluster node.

[Top](#)

On x86 Systems the CPU Power Management Feature Is Disabled During Cluster Configuration (6870631)

Problem Summary: On all platforms, the cluster configuration process disables power management by renaming the file `/etc/power.conf` to `/etc/power.conf.date_stamp`. However, on systems based on x86 processors running the Solaris 10 10/09 OS or the equivalent patch, the power management feature does not interfere with cluster functioning and can remain on.

Workaround: After you configure an x86 based cluster, to reenble the CPU power management feature, perform the following steps:

1. Rename `/etc/power.conf.date_stamp` to `/etc/power.conf`.
2. Execute the `pmconfig` command for the change to take effect.

[Top](#)

The Telemetry Resource Fails to Come Online in the Solaris 9 OS When It Is Configured Using `clsetup` (6873451)

Problem Summary: On the Solaris 9 OS, when telemetry is configured by using the `clsetup` utility, telemetry fails because the telemetry resource does not come online. The following error messages appear:

```

/usr/cluster/lib/rgm/rt/sctelemetry/sctelemetry -i -o
hasp_mnt_pt=/global/mysql,hasp_rg=mysql-rg,hasp_rs=sql-stor,db_rg=cl-db-rg,db_rs=cl-d

Aug 19 16:07:40 pjoker2 genunix: NOTICE: core_log: sleep[5347] core dumped:
/var/core/core.sleep.5347.1250678260.0
Aug 19 16:07:40 pjoker2 genunix: NOTICE: core_log: sctelemetry_svc_[4987] core
dumped: /var/core/core.sctelemetry_svc_.4987.1250678260.0
Aug 19 16:07:40 pjoker2 Cluster.rgmd: Method <sctelemetry_svc_start> on resource
<cl-tlmtry-rs>, resource group <cl-tlmtry-rg>, node <pjoker2>: Timeout.
Error: bringing resource group "cl-tlmtry-rg" online failed
scswitch: (C450598) On node pjoker1, resource group cl-tlmtry-rg is online but
resource cl-tlmtry-rs failed to start
(C450598) On node pjoker2, resource group cl-tlmtry-rg is online but resource
cl-tlmtry-rs failed to start
(C450598) On node pjoker3, resource group cl-tlmtry-rg is online but resource
cl-tlmtry-rs failed to start
(C450598) On node pjoker4, resource group cl-tlmtry-rg is online but resource
cl-tlmtry-rs failed to start

Command failed.

```

Workaround: Disable IPv6 by removing from the `/etc/inet/ipnodes` file the entry that starts with `::1`.

[Top](#)

The Command `clnode remove -F nodename` Fails to Remove the Node `nodename` From Solaris Volume Manager Device Groups (6471834)

Problem Summary: When a node is removed from the cluster by using the command `clnode remove -F nodename`, a stale entry for the removed node might remain in Solaris Volume Manager device groups.

Workaround: Remove the node from the Solaris Volume Manager device group by using the `metaset` command before you run the `clnode remove -F nodename` command.

If you ran the `clnode remove -F nodename` command before you removed the node from the Solaris Volume Manager device group, run the `metaset` command from an active cluster node to remove the stale node entry from the Solaris Volume Manager device group. Then run the `clnode clear -F nodename` command to completely remove all traces of the node from the cluster.

[Top](#)

If the Solaris Security Toolkit Is Configured on Cluster Nodes, `scstat -i` Gives an RPC Bind Failure Error(6727594)

Problem Summary: If the Solaris Security Toolkit is configured on cluster nodes, the command `scstat -i` gives an RPC bind failure error. The error message is similar to the following:

```

scrconf: RPC: Rpcbnd failure - RPC: Authentication error

```

Other Sun Cluster commands that use RPC, such as `clsnmpuser`, might also fail.

Workaround: Add the cluster private hostnames or the IP addresses associated with the cluster private hostnames to the `/etc/hosts.allow` file.

[Top](#)

Configuring a Scalable Resource With the `LB_STICKY_WILD` Load Balancing Policy Fails With `clsetup` (6773401)

Problem Summary: You cannot use the `clsetup` utility to configure a resource with the load balancing policy `LB_STICKY_WILD`. The policy is set to `LB_WILD` instead.

Workaround: Use the `clresource create` command to set the load balancing policy to `LB_STICKY_WILD`.

[Top](#)

The Sun Cluster Networking SSM Layer Must Be Zone-Cluster Aware (6885244)

Problem Summary: If a scalable resource that is configured in the global cluster or in a zone cluster shares the same name with another scalable resource that is configured in a different zone cluster, a change in configuration of the load balancer in one zone cluster will change the load balancer configuration in the other zone cluster as well. This behavior breaks the delegated administration model of zone clusters.

Workaround: Configure scalable resources with unique names across the global cluster and all zone clusters.

[Top](#)

iSCSI-Based Devices Do Not Work in Clusters on the Solaris 10 10/09 OS (6888193)

Problem Summary: On the Solaris 10 10/09 OS, the boot order of the iSCSI initiator changed, as a bug fix for CR 6786839. On a cluster, this change causes the iSCSI initiator to start after the DID drivers are loaded. As a result, functionality that is based on having these DID drivers be present at boot does not work. For iSCSI devices, you must manually load the DID drivers after the boot.

This behavior affects the following cluster components:

- **Quorum - Sun Cluster 3.2 11/09 software supports iSCSI based quorum on Sun StorageTek 2510 RAID arrays and Sun Storage 7000 series storage systems.** Quorum devices that are configured by using these devices through iSCSI do not contribute a quorum vote to a node at boot time. Therefore, cluster nodes might fail to boot, because there is no available quorum device.
- **iSCSI based storage support - Sun Cluster 3.2 11/09 software supports the Sun Storage 7000 series of systems as iSCSI based block devices.** Devices that are configured this way do not have the DID drivers loaded at boot. Any application that relies on having the DID drivers present at boot for these devices does not work.

Workaround: There is no workaround for the problem during boot. After the node boot is complete, manually load the DID drivers for the iSCSI based devices by using the `scdidadm -ui` command. For example, the following command loads the DID drivers for all the DID devices that are attached to a single node:

```
# /usr/cluster/bin/scdidadm -ui
```

Run this command on each node of the cluster.

[Top](#)

cl_apid Core Dump clapi and clapi_fi tests Fails in Solaris (6889682)

Problem Summary: On the Solaris 9 OS, if a resource is configured with the resource type `SUNW.event`, which is an implementation of Cluster Reconfiguration Notification Protocol (CRNP), the `cl_apid` daemon dumps core and does not deliver any event notifications.

Workaround: Avoid using CRNP notifications, or do not upgrade to Sun Cluster 3.2 11/09 software on the Solaris 9 OS.

[Top](#)

During HAStoragePlus Testing, rgmd Died and Failfast: Aborting zone "global" (6888345)

Problem Summary: On the Solaris 10 OS, if a zone-cluster node is in the process of booting into cluster mode and simultaneously a zone-cluster resource or resource group is operated upon using command line interfaces such as `clresourcegroup` or `clresource`, then an `rgmd` core dump might cause a global-cluster node to reboot. This problem occurs when at least one of

the zone-cluster resources has a method that has to be executed on the zone-cluster node which is still in the process of booting up.

Workaround: Wait for the zone-cluster node to boot into cluster mode before executing any of the commands. This can be verified using the output of 'clzonecluster status' command where the 'Status' for a zone-cluster node will be reported as 'Online' if it has successfully booted into cluster mode.

[Top](#)

Autodiscovery Does Not Work on LDomS With Hybrid I/O (6870171)

Problem Summary: During cluster configuration on LDomS with hybrid I/O, autodiscovery does not report any paths for the cluster interconnect.

Workaround: When you run the interactive `scinstall` utility, choose to configure the sponsor node and additional nodes in separate operations, rather than by configuring all nodes in a single operation. When the utility prompts "Do you want to use autodiscovery?", answer "no". You can then select transport adapters from the list that is provided by the `scinstall` utility.

[Top](#)

Misleading Error Message Output When Resource Type Must Be Removed From Previously-Removed Node Prior To Running `clnode clear` (6834733)

Problem Summary: After the `clnode remove` command removes a node from the cluster, the `clnode clear -F` command might fail if a resource type that is still registered on another cluster node was also registered on the removed node. The error message looks similar to the following:

```
clnode: (C109852) SUNW.HAStoragePlus is not a valid zone cluster
```

Workaround: The error message itself does not affect cluster function. To resolve the issue, remove the previously-removed node from those nodes on which the resource type is registered before you run the `clnode clear` command. Use the following syntax:

```
# clresourcetype remove-node -n nodename resource_type
```

You must specify the resource type without a version number.

[Top](#)

Non-Interactive `scinstall` Without the `-w` Option Sets the Wrong Cluster Netmask (6879147)

Problem Summary: If the non-interactive `scinstall` command is used without the `-w` option to install the Sun Cluster software, it sets a default cluster netmask that is not large enough to accommodate twelve zone clusters, the default number. When the user attempts to create a zone cluster, the zone cluster creation command `clzonecluster` might fail with an error message similar to the following:

```
No subnet available
Failed to assign a subnet for zone cmszone
```

Workaround: Use the following command to update the cluster netmask, to create enough subnets for the specified number of zone clusters. Run the command on only one cluster node:

```
# cluster set-netprops -p num_zoneclusters=number_of_zone_clusters
```

[Top](#)

Data Services

Validation Should Fail for a Scalable Resource With the `Outgoing_Connection` Property Set and With Multiple Entry Points (6886818)

Problem Summary: The problem occurs when the `Outgoing_Connection` property is set on a scalable resource and the resource is configured with multiple failover resource groups, each containing one or more scalable addresses. Such a configuration is not caught during validation, but TCP connections for the scalable service might be disrupted as a result.

Workaround: If the `Outgoing_Connection` property is enabled for a scalable resource, include all its scalable addresses in a single failover resource group.

[Top](#)

Shared Address Wizard Should Include Zone Cluster Nodes in the List of Available Nodes and Zones(6873218)

Problem Summary: Data-service configuration wizards in the `clsetup` utility and in Sun Cluster Manager do not configure applications in zone clusters. Only the Oracle RAC wizard is currently zone cluster aware. All others wizards only configure applications in the global cluster.

Workaround: Use the Sun Cluster command-line interface to configure applications in zone clusters.

[Top](#)

Installation

Autodiscovery Does Not Work for the `qlge` Driver for PCIe FCoE CNA (6939847)

Problem Summary: During an Oracle Solaris Cluster installation, auto discovery for the `qlge` driver for the PCIe LP and PCIe ExpressModule FCoE Converged Network Adapters (CNAs) does not work. The following products are affected:

- Oracle's Sun Storage 10GbE FCoE PCIe CNA
- Oracle's Sun Storage 10GbE FCoE ExpressModule CNA

Workaround: When you run the `scinstall` utility and you are prompted for the interconnect adapters, select Other and type the name of each `qlge` interface.

[Top](#)

Autodiscovery Should Find Only One Interconnect Path for Each Adapter (6299097)

Problem Summary: If there are redundant paths in the network hardware between interconnect adapters, the `scinstall` utility might fail to configure the interconnect path between them.

Workaround: If autodiscovery discovers multiple interconnect paths, manually specify the adapter pairs for each path.

[Top](#)

The `installer` Utility Installs Solaris 8 Instead of Solaris 9 Shared Components During Sun Cluster Live Upgrade (6870341)

Problem Summary: A live upgrade from Sun Cluster 3.1 software on the Solaris 8 OS to Sun Cluster 3.2 11/09 software on the Solaris 9 9/05 OS loads the wrong version of the shared components to the alternate root, so the upgrade fails.

Workaround: Run the `installer` utility twice to put the correct version of shared components in the alternate root.

[Top](#)

The `installer` Utility Hangs in Solaris 10 Exclusive-IP Zone Configuration (6871624)

Problem Summary: On the Solaris 10 OS, the `installer` utility hangs if you create a non-global zone and configure the zone's IP type as `exclusive` but you do not assign the non-global zone an IP address.

Workaround: Include an IP address when you create an exclusive-IP non-global zone by using the `installer` utility.

[Top](#)

Node Hangs On Reboot After Application of KU 141444-08 or Later Patch (6898395)

Problem Summary: After KU 141444-08 or later versions of the patch are applied to a node, the node cannot rejoin the cluster because all autopush structures are exhausted. The boot process hangs after these messages are displayed on the console:

```
Copyright 1983-2009 Sun Microsystems, Inc. All rights reserved.
Use is subject to license terms.
```

Workaround: Add the following parameter to the `/etc/system` file on all cluster nodes before rebooting:

```
set nautopush=64
```

[Top](#)

Localization

[zh_CN]: Result of System Requirements Checking Is Wrong (6495984)

Problem Summary: When you use the `installer` utility in the Simplified Chinese and Traditional Chinese locales to install Sun Cluster software, the software that checks the system requirements incorrectly reports that the swap space is 0 Mbytes.

Workaround: Ignore this reported information. In these locales, you can run the following command to determine the correct swap space:

```
# df -h | grep swap
```

[Top](#)

Runtime

HAStoragePlus Resources Fail To Import Using Cachefile After Kernel Patch 141444-09 Is Applied (6895580)

Problem Summary: If Solaris patch 141444-09 (SPARC) or 141445-09 (x86) is applied, **HAStoragePlus** resources will import ZFS zpools without using the `zpool.cachefile`, slowing the import. If there are many zpools or zpools with many disks, the import might take hours.

Workaround: Apply the latest version of the Sun Cluster 3.2 CORE patch: Patch 126106 for Solaris 10 SPARC ,or patch 126107 for Solaris 10 x86.

[Top](#)

Performance Degradation of More Than 10% of Net Reconfiguration Time (6871753)

Problem Summary: On a heavily loaded system with Solaris Volume Manager for Sun Cluster multi-owner disksets, it might take more than 10% longer to complete a cluster reconfiguration cycle on the Solaris 10 10/09 OS than it did on the Solaris 10 10/08 OS.

Workaround: There is currently no workaround. Contact your Sun service representative to determine whether a patch is available.

[Top](#)

SMF Service Start Method for `sc_synicsa_server` Hangs After Boot (6896430)

Problem Summary: Even if no scalable service is configured to use the IPsec session failover feature, during the boot process the new SMF service `svc:/system/cluster/sc_synicsa_server`, which enables IPsec session failover, attempts to come online, whereupon the `sc_synicsa_server` service start method fails to complete and hangs:

```
# svcs -x
...
svc:/system/cluster/sc_synicsa_server:default (SA synchronization service)
State: offline since Thu Oct 29 09:04:23 2009
Reason: Start method is running.
```

This failure might impact other SMF services, including those services that are required for cluster operations such as `svc:/system/cluster/scdpm` and `svc:/system/cluster/pnm`.

Workaround:

Manually kill the `sc_synicsa_server` start method and the `synicsa_serverd` process. After you have killed the `sc_synicsa_server` start method and the `synicsa_serverd` process, then:

- If the scalable service feature for IPsec Session Failover is used, run the following command:

```
# svcadm enable sc_synicsa_server
```

- If the scalable service feature for IPsec Session Failover is not used, run the following command:

```
# svcadm disable sc_synicsa_server
```

Alternately, if the scalable service feature for IPsec Session Failover is not used and you cannot manually kill the `sc_synicsa_server` start method and the `synicsa_serverd` process, run the following command:

```
# svcadm disable sc_synicsa_server
```

Then kill the `sc_synicsa_server` start method and the `synicsa_serverd` process.

[Top](#)

Upgrade

In Rolling Upgrade Scenario, the Command `scversions` Hangs If It Is Run Before New Software Is Installed On All Nodes (6874639)

Problem Summary: In a rolling upgrade, the `scversions` command might hang if not all of the nodes are installed with the latest software.

Workaround: Do not run the `scversions` command until all the nodes are installed with the new software. To check whether new software has been installed on a node without using `scversions`, see the `/etc/cluster/release` file on the node, which lists the last installed or upgraded version of Sun Cluster software.

[Top](#)

The Command `zoneadm` Should Mount `cluster` Branded Zones As `native` in

Alternate Roots: Solaris Upgrade Fails in Presence of Zone Clusters (6874636)

Problem Summary: If a zone cluster is configured on a global-cluster node that runs the Solaris 10 OS, the use of JumpStart to upgrade the OS of that global-cluster node fails with a warning message like the following:

```
Warning
The Solaris Version (Solaris 10) on slice c1t1d0s0 cannot be upgraded.

A non-global zone could not be mounted.
```

The problem affects all JumpStart upgrades of the Solaris 10 OS to at least the Solaris 10 10/08 OS, if there are zone clusters installed on the global-cluster node. For example, an upgrade from the Solaris 10 5/08 OS to the Solaris 10 10/08 OS fails, as does an upgrade from the Solaris 10 5/09 OS to the Solaris 10 10/09 OS.

Note:

- (1) The problem does not occur if no zone clusters are installed on the cluster being upgraded.
- (2) The problem does not occur if there are zone clusters, but they exist only in the `Configured` state.

Workaround: Add two files to the Solaris installation miniroot image on the JumpStart server, as shown in the steps below. In these instructions, `server` is the name of the JumpStart server that you use to install the upgrade. Perform all steps as `superuser`.

Perform Step 1 from one node of the cluster that is to be upgraded.

1. Copy the zone-cluster `config.xml` and `platform.xml` files from a node that is to be upgraded to a location on the JumpStart server.

```
# cp /usr/lib/brand/cluster/config.xml.upgrade /net/server/some_dir/config.xml
# cp /usr/lib/brand/cluster/platform.xml /net/server/some_dir/platform.xml
```

Perform Steps 2-6 on the JumpStart server.

2. Create a new Solaris installation image.

```
# cd path_to_Solaris_install_image/Solaris_10/Tools
# ./setup_install_server path_to_new_Solaris_install_image
# cd
```

3. Unpack the miniroot.

```
# /boot/solaris/bin/root_archive unpackmedia path_to_new_Solaris_install_image
path_to_miniroot
```

4. Put the zone-cluster brand files in the miniroot.

```
# mkdir path_to_miniroot/usr/lib/brand/cluster
# cp /net/server/some_dir/config.xml path_to_miniroot/usr/lib/brand/cluster
# cp /net/server/some_dir/platform.xml path_to_miniroot/usr/lib/brand/cluster
```

5. Pack the miniroot and place it in the new installation image.

```
# /boot/solaris/bin/root_archive packmedia path_to_new_Solaris_install_image
path_to_miniroot
```

6. Run the `addclient` command from the new installation image location for each cluster node to be upgraded.

```
# cd path_to_new_Solaris_install_image/Solaris_10/Tools
# ./addclient your_configuration_details
```

Perform this workaround just once for each JumpStart server and for each Solaris update to which you plan to upgrade.

[Top](#)

Sun Cluster Upgrade with `scinstall -u update` Fails When Configured Zones Are Not Booted (6887976)

Problem Summary: Upgrading Sun Cluster software by using the command `scinstall -u update` fails when configured non-global zones are not booted. Package removal fails with error messages similar to the following:

```
Removal of package-name on zone zone_name failed (internal error).
No changes were made to the system.
```

Workaround: Ensure that all `native` and `cluster` brand non-global zones on the global-cluster node that you upgrade are booted before you run the `scinstall -u update` command.

If some non-global zones are down so that the `scinstall -u update` command aborts, and some packages have been removed from the global zone, do the following:

1. Boot all `native` and `cluster` brand zones.

```
# cd /usr/lib/brand/cluster
# cp config.xml config.xml.upgrade
# cp DVD/Solaris_sparc/Product/sun_cluster/Solaris_n
/Packages/SUNWscu/reloc/usr/lib/brand/cluster/platform.xml platform.xml
```

2. Run the `scinstall -u update` command again.

[Top](#)

Setting the `num_zoneclusters` Property Causes a Core Dump on a Cluster That Was Upgraded from Sun Cluster 3.1 to Sun Cluster 3.2 1/09 Software (6846142)

Problem Summary: A new property `num_zoneclusters` that was introduced in Sun Cluster 3.2 1/09 software defines the maximum number of zone clusters that can be created on the cluster. When a cluster is upgraded to Sun Cluster 3.2 1/09 software from a version earlier than the Sun Cluster 3.2 1/09 release, you must set this property by using the `cluster` command before you configure any zone clusters. However, if the cluster is upgraded to Sun Cluster 3.2 11/09 from Sun Cluster 3.1 software, the process of setting the `num_zoneclusters` property generates a core dump.

Workaround: Set the properties `max_nodes` and `max_privatenets` that were introduced in the initial Sun Cluster 3.2 release, before you set the `num_zoneclusters` property. Use the following procedure:

- 1) Reboot all nodes into non-cluster mode.
- 2) Use the `/usr/cluster/bin/clsetup` or `/usr/cluster/bin/cluster` commands to set the values of the properties `max_nodes` and `max_privatenets`.
- 3) Reboot all nodes into cluster mode.
- 4) Run the `/usr/cluster/bin/cluster` command to set the value of `num_zoneclusters`.

For the command line syntax use:

```
# cluster set-netprops -p private_netaddr="172.16.0.0" -p max_nodes=
number_of_nodes_expected_in_the_cluster -p max_privatenets=
number_of_networks_expected_in_the_cluster
```

In the `clsetup` utility make selections similar to those shown:

```

*** Main Menu ***

1) Change Network Addressing and Ranges for the Cluster Transport
2) Show Network Addressing and Ranges for the Cluster Transport

?) Help with menu options
q) Quit

___ Option: 1

>>> Change Network Addressing and Ranges for the Cluster Transport <<<

Network addressing for the cluster transport is currently configured
as follows:

Private Network ===
private_netaddr: 172.16.0.0
___ private_netmask: 255.255.240.0
___ max_nodes: 64
___ max_privatenets: 10
___ num_zoneclusters: 12

Do you want to change this configuration (yes/no) [yes]? yes

The default network address for the cluster transport is 172.16.0.0.

Do you want to use the default (yes/no) [yes]? yes

The default netmask for the cluster transport is 255.255.240.0.

Do you want to use the default (yes/no) [yes]? no

The combination of private netmask and network address will dictate
both the maximum number of nodes and private networks that can be
supported by a cluster. Given your private network address, this
program will generate a range of recommended private netmasks based on
the maximum number of nodes and private networks that you anticipate
for this cluster.

In specifying the anticipated number of maximum nodes and private
networks for this cluster, it is important that you give serious
consideration to future growth potential. While both the private
netmask and network address can be changed later, the tools for making
such changes require that all nodes in the cluster be booted into
noncluster mode.

Maximum number of nodes anticipated for future growth [64]? 64

Maximum number of private networks anticipated for future growth [10]? 10

Specify a netmask of 255.255.248.0 to meet anticipated future
requirements of 64 cluster nodes and 10 private networks.

To accommodate more growth, specify a netmask of 255.255.240.0 to
support up to 64 cluster nodes and 20 private networks.

What netmask do you want to use [255.255.248.0]? 255.255.248.0
Is it okay to proceed with the update (yes/no) [yes]? yes

```

Top

Zones With ip-type=exclusive Cannot Host SUNW.LogicalHostname Resources After Upgrade (6702621)

Problem Summary: Resource type `SUNW.LogicalHostname` is registered at version 2 (use the `clresourcetype list` command to display the version). After upgrade, logical-hostname resources can be created for zones with `ip-type=exclusive`, but network access to the logical hostname, for example, telnet or rsh, does not work.

Workaround: Perform the following steps:

1. Delete all resource groups with a node list that contains a zone with `ip-type=exclusive` that hosts logical-hostname resources.
2. Upgrade the `SUNW.LogicalHostname` resource type to version 3:

```
# clresourcetype register SUNW.LogicalHostname:3
```

[Top](#)

After Live Upgrade Some Cluster Services Remain in the Disabled State (6898156)

Problem Summary: After a live upgrade from Sun Cluster 3.2 GA to Sun Cluster 3.2 11/09 software, some of the cluster services fail to come up and remain in the disabled state. Error messages similar to the following are displayed on the console:

```
Rebooting with command: boot
Boot device: /pci@1c,600000/scsi@2/disk@1,0:a File and args: kmdb
Loading kmdb...
SunOS Release 5.10 Version Generic_141444-09 64-bit
Copyright 1983-2009 Sun Microsystems, Inc. All rights reserved.
Use is subject to license terms.
WARNING: todm5819p_rmc: kernel debugger detected: hardware watchdog disabled
Hostname: pfear3
Nov 3 13:42:28 svc.startd[8]: svc:/application/management/snmpdx:default: Method
"/lib/svc/method/svc-snmpdx" failed with exit status 95.
Nov 3 13:42:28 svc.startd[8]: application/management/snmpdx:default failed fatally:
transitioned to maintenance (see 'svcs -xv' for details)
Configuring devices.
Nov 3 13:42:44 in.mpathd[160]: missed sending 2 probes cur_time 92596 snxt_time
93030 snxt_basetime 91664
Nov 3 13:42:44 in.mpathd[160]: missed sending 1 probes cur_time 92599 snxt_time
92661 snxt_basetime 91263
NIS domain name is scdev.sfbay.sun.com
Nov 3 13:42:54 svc.startd[8]: svc:/application/print/cleanup:default: Method
"/lib/svc/method/print-cleanup start" failed with exit status 1.
Nov 3 13:42:54 svc.startd[8]: svc:/application/print/cleanup:default: Method
"/lib/svc/method/print-cleanup start" failed with exit status 1.
Nov 3 13:42:54 svc.startd[8]: svc:/application/print/cleanup:default: Method
"/lib/svc/method/print-cleanup start" failed with exit status 1.
Nov 3 13:42:54 svc.startd[8]: application/print/cleanup:default failed:
transitioned to maintenance (see 'svcs -xv' for details)
Loading smf(5) service descriptions: 15/114
Will not sync "/etc/cluster/ccr/global" from last BE - no such file or directory.
Will not sync "/etc/cluster/ccr_ql.bak" from last BE - no such file or directory.
Will not sync "/etc/cluster/ql_data.txt" from last BE - no such file or directory.
Will not sync "/etc/cluster/upgrade_partition_a_run" from last BE - no such file or
directory.
17/114
Cluster Configuration Repository transformation completed successfully
18/114
Booting in cluster mode
19/114
NOTICE: CMM: Node pfear1 (nodeid = 1) with votecount = 1 added.
NOTICE: CMM: Node pfear3 (nodeid = 2) with votecount = 1 added.
NOTICE: CMM: Node pfear4 (nodeid = 3) with votecount = 1 added.
NOTICE: CMM: Node pfear2 (nodeid = 4) with votecount = 1 added.
20/114
```

```
NOTICE: CMM: Quorum device 1 (/dev/did/rdisk/d3s2) added; votecount = 3, bitmask of
nodes with configured paths = 0xf.
21/114
NOTICE: clcomm: Adapter cel constructed
NOTICE: clcomm: Adapter ce0 constructed
NOTICE: CMM: Node pfear3: attempting to join cluster.
25/114
NOTICE: CMM: Node pfear1 (nodeid: 1, incarnation #: 1257284579) has become
reachable.
NOTICE: clcomm: Path pfear3:cel - pfear1:cel online
27/114
NOTICE: clcomm: Path pfear3:ce0 - pfear1:ce0 online
44/114
NOTICE: CMM: Node pfear4 (nodeid: 3, incarnation #: 1257284477) has become
reachable.
NOTICE: clcomm: Path pfear3:cel - pfear4:cel online
45/114
NOTICE: clcomm: Path pfear3:ce0 - pfear4:ce0 online
91/114
NOTICE: CMM: Node pfear2 (nodeid: 4, incarnation #: 1257284594) has become
reachable.
NOTICE: clcomm: Path pfear3:ce0 - pfear2:ce0 online
NOTICE: clcomm: Path pfear3:cel - pfear2:cel online
NOTICE: CMM: Cluster has reached quorum.
NOTICE: CMM: Node pfear1 (nodeid = 1) is up; new incarnation number = 1257284579.
NOTICE: CMM: Node pfear3 (nodeid = 2) is up; new incarnation number = 1257284580.
NOTICE: CMM: Node pfear4 (nodeid = 3) is up; new incarnation number = 1257284477.
NOTICE: CMM: Node pfear2 (nodeid = 4) is up; new incarnation number = 1257284594.
NOTICE: CMM: Cluster members: pfear1 pfear3 pfear4 pfear2.
93/114
NOTICE: CMM: node reconfiguration #3 completed.
NOTICE: CMM: Node pfear3: joined cluster.
NOTICE: CCR: Waiting for repository synchronization to finish.
102/114
Nov 3 13:43:22 Cluster.SMF.DR: Error in scha_cluster_get
114/114
ip: joining multicasts failed (18) on clprivnet0 - will use link layer broadcasts
for multicast
```

```
Configuring DID devices
Reading ZFS config: done.
obtaining access to all attached disks
```

Workaround: Manually enable the following cluster services and any other cluster services that are in the `disabled` state:

```
# svcadm enable sc_pmmd
# svcadm enable cl_execd
# svcadm enable cznetd
```


[Top](#)

Patches and Required Firmware Levels

This section provides information about patches for Sun Cluster configurations, including the following subsections:

- [Applying the Sun Cluster 3.2 11/09 Core Patch](#)
- [Removing the Sun Cluster 3.2 11/09 Core Patch](#)
- [Patch Management Tools](#)
- [My Oracle Support Online](#)


If you are upgrading to Sun Cluster 3.2 11/09 software, see the Sun Cluster Upgrade Guide for Solaris OS. Applying a Sun Cluster 3.2 11/09 Core patch does not provide the same result as upgrading the software to the Sun Cluster 3.2 11/09 release.

 **Note**
Read the patch `README` file before applying or removing any patch.

If you are using the rebooting patch (`node`) method to install the Sun Cluster core patch 126105 (S9/SPARC), 126106 (S10/SPARC), or 126107 (S19/x64), you must have the 125510-02, 125511-02, or 125512-02 core patch, respectively, installed before you can install the 126105, 126106, or 126107 patch. If you do not have the 125510-02, 125511-02, or 125512-02 patch installed and want to install 126105, 126106, or 126107, you must use the rebooting cluster method.

See the following list for examples of patching scenarios:

- If you have Sun Cluster 3.2 11/09 software using the Solaris 10 Operating System on SPARC with patch 125511-02 installed and want to install 126106, use the rebooting `node` or rebooting cluster method.
- If you have Sun Cluster 3.2 11/09 software using the Solaris 10 Operating System on SPARC without 125511-02 installed and want to install 126106, choose one of the following methods:
 - Use the rebooting cluster method to install 126106.
 - Install 125511-02 by using the rebooting `node` method and then install 126106 by using the rebooting `node` method.

 **Note**
You must be a registered My Oracle Support™ user to view and download the required patches for the Sun Cluster product. If you do not have a My Oracle Support account, contact your Oracle service representative or sales engineer, or register online at [My Oracle Support](#).

[Top](#)

Applying the Sun Cluster 3.2 11/09 Core Patch

Complete the following procedure to apply the Sun Cluster 3.2 11/09 core patch. Ensure that all nodes of the cluster are maintained at the same patch level.

How to Apply the Sun Cluster 3.2 11/09 Core Patch

1. Install the patch by 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. Perform Resource Type upgrades to any new versions of the resource type available. Run `clsetup` to obtain the list of new resource types. For more details, see the documentation on "Upgrading a Resource Type" in the Sun Cluster Data Services Planning and Administration Guide for Solaris OS.

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



Caution

If the Sun Cluster 3.2 11/09 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 requires planned downtime of these services. Therefore, do not perform Step 3 until you are ready to commit the Sun Cluster 3.2 11/09 core patch permanently to your cluster.

[Top](#)

Removing the Sun Cluster 3.2 11/09 Core Patch

Complete the following procedure to remove the Sun Cluster 3.2 11/09 core patch.

How to Remove the Sun Cluster 3.2 11/09 Core Patch

1. List the resource types on the cluster.

```
# clresourcetype list
```

2. If you have upgraded to any new resource types after applying the core patch then follow the directions in "How to Remove a Resource Type" in the Sun Cluster Data Services Planning and Administration Guide for Solaris OS.

1. Remove the Sun Cluster 3.2 11/09 core patch from each node on which you installed the patch.

```
# patchrm patch-id
```

2. Reboot into cluster mode all of the nodes from which you removed the Sun Cluster 3.2 11/09 core patch. Rebooting all of the nodes from which you removed the Sun Cluster 3.2 11/09 core patch before rebooting any unaffected nodes ensures that the cluster is formed with the correct configuration information on all nodes. If all nodes on the cluster were patched with the core patch, you can reboot the nodes into cluster mode in any order.
3. Reboot any remaining nodes into cluster mode.

For instructions about rebooting nodes into cluster mode, see "How to Reboot a Cluster Node" in the Sun Cluster System Administration Guide for Solaris OS.

[Top](#)

Patch Management Tools

Information about patch management options for the Solaris OS is available at the web sites for [Sun Ops Center](#) and [Sun Connection Update Manager](#).

The following tools are part of the Solaris OS. Refer to the version of the manual that is published for the Solaris OS release that is installed on your system:

- Information for using the Solaris patch management utility, `patchadd`, is provided in Solaris Administration Guide: Basic Administration at <http://docs.sun.com>.
- Information for using Solaris Live Upgrade to apply patches is provided in the Solaris installation guide for Live Upgrade and upgrade planning at <http://docs.sun.com>.

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.

[Top](#)

Patch for Cluster Support for Sun StorageTek 2530 Array

The Sun StorageTek Common Array Manager (CAM) software, Version 6.0.1, provides SCSI3 or PGR support for the Sun StorageTek 2530 array for up to three nodes. The patch is not a required upgrade for the Sun StorEdge 6130, 2540, 6140, and 6540, and StorageTek FLX240, FLX280 and FLX380 platforms. The CAM 6.0.1 patch is available from the Sun Download Center.

[Top](#)

My Oracle Support Online

The My Oracle Support™ Online Web site provides 24-hour access to the most up-to-date information regarding patches, software, and firmware for Oracle products. Access the My Oracle Support Online site at [My Oracle Support](#) for the most current matrixes of supported software, firmware, and patch revisions.

Sun Cluster 3.2 11/09 third-party patch information is provided through an Oracle Support Info Doc (formerly 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 11/09 environment. To locate this Info Doc, log on to My Oracle Support. From the My Oracle Support home page, type **Sun Cluster 3.x Third-Party Patches** in the search criteria box.

Before you install Sun Cluster 3.2 11/09 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 "[Patching Sun Cluster Software and Firmware](#)" in Sun Cluster System Administration Guide for Solaris OS.

[Top](#)

Sun Cluster 3.2 11/09 Documentation

The Sun Cluster 3.2 11/09 user documentation set consists of the following collections:

- [Sun Cluster 3.2 11/09 Software Manuals for Solaris OS](#)
- [Sun Cluster 3.2 11/09 Reference Manuals for Solaris OS](#)
- [Sun Cluster 3.2 11/09 Data Service Manuals for Solaris OS \(SPARC Platform Edition\)](#)
- [Sun Cluster 3.2 11/09 Data Service Manuals for Solaris OS \(x86 Platform Edition\)](#)
- [Sun Cluster 3.1-3.2 Hardware Collection for Solaris OS \(SPARC Platform Edition\)](#)
- [Sun Cluster 3.1-3.2 Hardware Collection for Solaris OS \(x86 Platform Edition\)](#)



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 [Sun Cluster 3.1 Data Service for Sun ONE Directory Server](#).

Beginning with Version 5.2 of Sun ONE Directory Server, see the [Sun ONE Directory Server](#) or [Sun Java System Directory Server](#) installation documentation.

Some manuals were not updated since the previous Sun Cluster 3.2 release. Their content, however, also applies to the Sun Cluster 3.2 11/09 release.

The Sun Cluster 3.2 11/09 user documentation is available in PDF and HTML format at the following web site:

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



Note

The Sun Cluster Data Service for Oracle Guide is translated into Japanese and Simplified Chinese.

Documentation for other data services, including man pages and wizard online help, is available only in English.

[Top](#)

Sun Cluster 3.2 11/09 Software Manuals for Solaris OS

Part Number	Book Title
820-7360	Sun Cluster 3.2 11/09 Documentation Center
821-0259	Sun Cluster Concepts Guide for Solaris OS
821-0261	Sun Cluster Data Services Developer's Guide for Solaris OS
820-7359	Sun Cluster Data Services Planning and Administration Guide for Solaris OS
820-7363	Sun Cluster Error Messages Guide for Solaris OS
821-0518	Sun Cluster Overview for Solaris OS
819-6811	Sun Cluster Quick Reference
820-4989	Sun Cluster Quick Start Guide for Solaris OS
820-7356	Sun Cluster Software Installation Guide for Solaris OS
820-7358	Sun Cluster System Administration Guide for Solaris OS
820-7357	Sun Cluster Upgrade Guide for Solaris OS

[Top](#)

Sun Cluster 3.2 11/09 Reference Manuals for Solaris OS

Part Number	Book Title
820-7361	Sun Cluster Reference Manual for Solaris OS
820-7362	Sun Cluster Data Services Reference Manual for Solaris OS
820-4686	Sun Cluster Quorum Server Reference Manual for Solaris OS

[Top](#)

Sun Cluster 3.2 11/09 Data Service Manuals for Solaris OS (SPARC Platform Edition)

Part Number	Book Title
820-5026	Sun Cluster Data Service for Agfa IMPAX Guide for Solaris OS
820-7092	Sun Cluster Data Service for Apache Guide for Solaris OS
819-3057	Sun Cluster Data Service for Apache Tomcat Guide for Solaris OS
819-3058	Sun Cluster Data Service for DHCP Guide for Solaris OS
819-2977	Sun Cluster Data Service for DNS Guide for Solaris OS
820-5024	Sun Cluster Data Service for Informix Guide for Solaris OS
819-5415	Sun Cluster Data Service for Kerberos Guide for Solaris OS
820-5034	Sun Cluster Data Service for MaxDB Guide for Solaris OS
821-0273	Sun Cluster Data Service for MySQL Guide for Solaris OS
819-3060	Sun Cluster Data Service for N1 Grid Service Provisioning System for Solaris OS

820-2565	Sun Cluster Data Service for NFS Guide for Solaris OS
821-0274	Sun Cluster Data Service for Oracle Guide for Solaris OS
820-2572	Sun Cluster Data Service for Oracle Application Server Guide for Solaris OS
821-0991	Sun Cluster Data Service for Oracle E-Business Suite Guide for Solaris OS
821-0262	Sun Cluster Data Service for Oracle RAC Guide for Solaris OS
820-5074	Sun Cluster Data Service for PostgreSQL Guide for Solaris OS
820-3040	Sun Cluster Data Service for Samba Guide for Solaris OS
820-5033	Sun Cluster Data Service for SAP Guide for Solaris OS
820-5035	Sun Cluster Data Service for SAP liveCache Guide for Solaris OS
820-2568	Sun Cluster Data Service for SAP Web Application Server Guide for Solaris OS
820-5036	Sun Cluster Data Service for Siebel Guide for Solaris OS
820-5025	Sun Cluster Data Service for Solaris Containers Guide
820-3042	Sun Cluster Data Service for Sun Grid Engine Guide for Solaris OS
820-5032	Sun Cluster Data Service for Sun Java System Application Server EE (HADB) Guide for Solaris OS
820-5029	Sun Cluster Data Service for Sun Java System Application Server Guide for Solaris OS
820-5031	Sun Cluster Data Service for Sun Java System Message Queue Guide for Solaris OS
820-5030	Sun Cluster Data Service for Sun Java System Web Server Guide for Solaris OS
821-0315	Sun Cluster Data Service for SWIFTAlliance Access Guide for Solaris OS
821-0314	Sun Cluster Data Service for SWIFTAlliance Gateway Guide for Solaris OS
820-2570	Sun Cluster Data Service for Sybase ASE Guide for Solaris OS
821-0272	Sun Cluster Data Service for WebLogic Server Guide for Solaris OS
819-3068	Sun Cluster Data Service for WebSphere Message Broker Guide for Solaris OS
819-3067	Sun Cluster Data Service for WebSphere MQ Guide for Solaris OS

[Top](#)

Sun Cluster 3.2 11/09 Data Service Manuals for Solaris OS (x86 Platform Edition)

Part Number	Book Title
820-7092	Sun Cluster Data Service for Apache Guide for Solaris OS
819-3057	Sun Cluster Data Service for Apache Tomcat Guide for Solaris OS
819-3058	Sun Cluster Data Service for DHCP Guide for Solaris OS
819-2977	Sun Cluster Data Service for DNS Guide for Solaris OS
820-5024	Sun Cluster Data Service for Informix Guide for Solaris OS
819-5415	Sun Cluster Data Service for Kerberos Guide for Solaris OS
820-5034	Sun Cluster Data Service for MaxDB Guide for Solaris OS
821-0273	Sun Cluster Data Service for MySQL Guide for Solaris OS
819-3060	Sun Cluster Data Service for N1 Grid Service Provisioning System for Solaris OS

820-2565	Sun Cluster Data Service for NFS Guide for Solaris OS
821-0274	Sun Cluster Data Service for Oracle Guide for Solaris OS
820-2572	Sun Cluster Data Service for Oracle Application Server Guide for Solaris OS
821-0262	Sun Cluster Data Service for Oracle RAC Guide for Solaris OS
820-5074	Sun Cluster Data Service for PostgreSQL Guide for Solaris OS
820-3040	Sun Cluster Data Service for Samba Guide for Solaris OS
820-5033	Sun Cluster Data Service for SAP Guide for Solaris OS
820-5035	Sun Cluster Data Service for SAP liveCache Guide for Solaris OS
820-2568	Sun Cluster Data Service for SAP Web Application Server Guide for Solaris OS
820-5025	Sun Cluster Data Service for Solaris Containers Guide
820-3042	Sun Cluster Data Service for Sun Grid Engine Guide for Solaris OS
820-5032	Sun Cluster Data Service for Sun Java System Application Server EE (HADB) Guide for Solaris OS
820-5029	Sun Cluster Data Service for Sun Java System Application Server Guide for Solaris OS
820-5031	Sun Cluster Data Service for Sun Java System Message Queue Guide for Solaris OS
820-5030	Sun Cluster Data Service for Sun Java System Web Server Guide for Solaris OS
820-2570	Sun Cluster Data Service for Sybase ASE Guide for Solaris OS
821-0272	Sun Cluster Data Service for WebLogic Server Guide for Solaris OS
819-3067	Sun Cluster Data Service for WebSphere MQ Guide for Solaris OS
819-3068	Sun Cluster Data Service for WebSphere Message Broker Guide for Solaris OS

[Top](#)

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

Part Number	Book Title
819-2993	Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS
819-2994	Sun Cluster 3.1 - 3.2 With Fibre Channel JBOD Storage Device Manual
819-3024	Sun Cluster 3.1 - 3.2 With Network-Attached Storage Devices Manual for Solaris OS
819-2995	Sun Cluster 3.1 - 3.2 With SCSI JBOD Storage Device Manual for Solaris OS
820-1692	Sun Cluster 3.1 - 3.2 With StorageTek 2540 RAID Arrays Manual for Solaris OS
819-7306	Sun Cluster 3.1 - 3.2 With StorageTek RAID Arrays Manual for Solaris OS
819-3015	Sun Cluster 3.1 - 3.2 With Sun StorEdge 3310 or 3320 SCSI RAID Array Manual for Solaris OS
819-3016	Sun Cluster 3.1 - 3.2 With Sun StorEdge 3510 or 3511 FC RAID Array Manual for Solaris OS
819-3017	Sun Cluster 3.1 - 3.2 With Sun StorEdge 3900 Series or Sun StorEdge 6900 Series System Manual
819-3018	Sun Cluster 3.1 - 3.2 With Sun StorEdge 6120 Array Manual for Solaris OS
819-3019	Sun Cluster 3.1 - 3.2 With Sun StorEdge 6130 Array Manual
819-3020	Sun Cluster 3.1 - 3.2 With Sun StorEdge 6320 System Manual for Solaris OS
819-3021	Sun Cluster 3.1 - 3.2 With Sun StorEdge 9900 Series Storage Device Manual for Solaris OS

819-2996	Sun Cluster 3.1 - 3.2 With StorEdge A1000 Array, Netra st A1000 Array, or StorEdge A3500 System Manual
819-3022	Sun Cluster 3.1 - 3.2 With Sun StorEdge A3500FC System Manual for Solaris OS
819-3023	Sun Cluster 3.1 - 3.2 With Sun StorEdge T3 or T3+ Array Manual for Solaris OS

[Top](#)

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

Part Number	Book Title
819-2993	Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS
819-3024	Sun Cluster 3.1 - 3.2 With Network-Attached Storage Devices Manual for Solaris OS
819-2995	Sun Cluster 3.1 - 3.2 With SCSI JBOD Storage Device Manual for Solaris OS
820-1692	Sun Cluster 3.1 - 3.2 With StorageTek 2540 RAID Arrays Manual for Solaris OS
819-7306	Sun Cluster 3.1 - 3.2 With StorageTek RAID Arrays Manual for Solaris OS
817-0180	Sun Cluster 3.1 - 3.2 With Sun StorEdge 3310 or 3320 SCSI RAID Array Manual for Solaris OS
819-3016	Sun Cluster 3.1 - 3.2 With Sun StorEdge 3510 or 3511 FC RAID Array Manual for Solaris OS
819-3018	Sun Cluster 3.1 - 3.2 With Sun StorEdge 6120 Array Manual for Solaris OS
819-3019	Sun Cluster 3.1 - 3.2 With Sun StorEdge 6130 Array Manual
819-3020	Sun Cluster 3.1 - 3.2 With Sun StorEdge 6320 System Manual for Solaris OS
819-3021	Sun Cluster 3.1 - 3.2 With Sun StorEdge 9900 Series Storage Device Manual for Solaris OS

[Top](#)

Localization

Localization for certain components of the Sun Cluster 3.2 11/09 software is as follows:

Component	Localization
Software Components	Japanese, Simplified Chinese
Online Help	Japanese, Simplified Chinese
Man Pages	Japanese

Past releases of Sun Cluster software included localization for languages that are no longer supported. In the Sun Cluster 3.2 11/09 release, some system messages and user interfaces might still be supplied with old translations in previously supported languages.

The following table shows the commands that set command line messages to English for commonly used shells:

shell	Command
sh	<code>\$ LC_MESSAGES=C;export LC_MESSAGES</code>
ksh	<code>\$ export LC_MESSAGES=C</code>
bash	<code>\$ export LC_MESSAGES=C</code>
csh	<code>% setenv LC_MESSAGES C</code>
tcsh	<code>% setenv LC_MESSAGES C</code>

[Top](#)

Documentation Issues

This section discusses errors or omissions for documentation, online help, or man pages in the Sun Cluster 3.2 11/09 release.

- [Software Installation Guide](#)
- [Data Services Planning and Administration Guide](#)
- [Data Service for Oracle Guide](#)
- [Data Service for Oracle Real Application Clusters Guide](#)
- [Man Pages](#)

[Top](#)

Software Installation Guide

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

Incorrect Restriction About NAS Devices for Quorum

Step 8 of "How to Configure Quorum Devices" contains the following note. This statement is incorrect and should be disregarded:



Note -
NAS devices are not a supported option for quorum devices in an Sun Cluster 3.2 11/09 configuration. Reference to NAS devices in the following table are for information only.

[Top](#)

Changes to Guidelines for Mediators

With the additional of support for three mediators, the guidelines listed in the "Configuring Dual-String Mediators" section are incorrect. Use the following new guidelines instead:

- A disk set can have up to three mediator hosts.
- The mediator host no longer needs to be part of the cluster.
- Mediators that are configured for disk sets must meet the existing two-string disk set criteria.
- The entire campus cluster can have more than two nodes.
- An N+1 cluster and other topologies are permitted.
- You must add the entry `root` to the `sysadmin` group in the `/etc/group` file on the mediator host. If the mediator host that you are adding does not yet have a disk set configured, you must also create a dummy disk set on that host. For example:

```
# metaset -s dummy-diskset-name -a -h nodename
```

To add the third mediator host, use the following `metaset` command to specify the additional node. This command replaces Step #2 of [How To Add Mediator Hosts](#) in the Sun Cluster Software Installation Guide for Solaris OS.



Note
Each node and disk set must have a unique name, and a disk set cannot be named `shared` or `admin`.

```
# metaset -s diskset-name -a -m clusternode1 clusternode2 othernode
```

[Top](#)

Data Services Planning and Administration Guide

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

Configuration Wizard Support for Zone Clusters Limited to Oracle RAC

The configuration of a data service in a zone cluster by using the Sun Cluster Manager or `clsetup` wizards is currently only supported for Oracle RAC. To configure other data services that are supported in a zone cluster, use the Sun Cluster maintenance commands.

[Top](#)

Data Service for Oracle Guide

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

Support for Third-Party Volume Managers with Clustered Oracle ASM Disk Groups

For Oracle 11g release 2, a patch and additional configuration procedures are required to support third-party volume managers with clustered Oracle ASM disk groups. Perform the following tasks to add this support:



Note

These instructions for HA for Oracle use procedures that are documented in the Sun Cluster Data Service for Oracle Real Application Clusters Guide.

Task	Additional Information
1. Apply at least the following minimum HA-Oracle patch version on each node that will run Oracle 11g release 2 with clustered Oracle ASM:	
<ul style="list-style-type: none"> • SPARC: 126047-16 • x86: 126048-16 	
2. Configure the disk set or disk group to use.	<p>Follow one of the following procedures in the Sun Cluster Data Service for Oracle Real Application Clusters Guide:</p> <ul style="list-style-type: none"> • If you are using Solaris Volume Manager for Sun Cluster, create a multi-owner disk set. Follow procedures in How to Create a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster for the Oracle RAC Database. • If you are using Veritas Volume Manager with the cluster feature, create a shared disk group. Follow procedures in How to Create a VxVM Shared-Disk Group for the Oracle RAC Database.
3. Perform procedures in Creating an Oracle Grid Infrastructure Resource .	When you configure Oracle RAC 11g release 2, this procedure replaces How to Create an Oracle Clusterware Resource for Interoperation With Sun Cluster in Sun Cluster Data Service for Oracle Real Application Clusters Guide.

[Top](#)

Creating an Oracle Grid Infrastructure Resource

Perform this procedure to manually create a Grid Infrastructure resource that proxies the Sun Cluster `SUNW.ScalDeviceGroup` resource. In this example procedure, the Grid Infrastructure resource is named `sun.resource`. You configure `sun.resource` to ensure that the corresponding Oracle ASM disk group is not mounted until `sun.resource` is online. The `sun.resource` resource comes online only if the corresponding `SUNW.ScalDeviceGroup` resource is online. And the `SUNW.ScalDeviceGroup` resource only comes online if the actual volume-manager disk set or disk group is online.

To ensure that the Oracle ASM disk group benefits from this dependency chain, after you define `sun.resource`, you modify the appropriate Grid Infrastructure Oracle ASM disk-group resource so that the hard-start dependency includes `sun.resource`. Modifying the hard-start dependency of the Grid Infrastructure Oracle ASM disk-group resource can only be performed by the `SUNW.scalable_asm_diskgroup_proxy` resource by using the `VALIDATE` method. Therefore, you must set an offline-restart dependency between the `SUNW.scalable_asm_diskgroup_proxy` and `SUNW.ScalDeviceGroup` resources.

[Top](#)

How to Create an Oracle Grid Infrastructure Resource

This example procedure sets the following dependencies:

- A Grid Infrastructure resource, `sun.resource`, where `resource` corresponds to a Sun Cluster `SUNW.ScalDeviceGroup` resource name:
 - o `sun.scal-asmdg1-rs`
 - o `sun.scal-asmdg2-rs`
- Sun Cluster offline-restart dependencies:
 - o `asm-data1-rs --> [add offline-restart dependency] --> scal-asmdg1-rs`
 - o `asm-data2-rs --> [add offline-restart dependency] --> scal-asmdg2-rs`

When an offline-restart dependency is set, the `VALIDATE` method of the `SUNW.scalable_asm_diskgroup_proxy` resource will transparently modify the appropriate Grid Infrastructure ASM disk group resource as follows:

- Grid Infrastructure hard-start dependencies:
 - o `ora.data1.dg --> [add hard-start dependency] --> sun.scal-asmdg1-rs`
 - o `ora.data1.dg --> [add hard-start dependency] --> sun.scal-asmdg1-rs`

Perform this procedure on one node of the cluster.

1. Become superuser.
2. Create the `sun.storage_proxy.type` resource type.

In this example procedure, the `/u01/app/11.2.0/grid/` directory is the Grid Infrastructure Oracle Home.

```
bash-3.00# /u01/app/11.2.0/grid/bin/crsctl add type sun.storage_proxy.type -basetype cluster_resource
```

3. Create `sun.resource` of type `sun.storage_proxy.type`.

[Top](#)

Removing an Oracle Grid Infrastructure Resource

This section contains the following procedures to remove an Oracle Grid Infrastructure resource.

- [How to Remove a Dependency](#)
- [How to Delete the `sun.resource` Resource](#)

[Top](#)

How to Remove a Dependency

This procedure shows how to set the offline-restart dependency to remove a dependency.

1. Display the current dependency.

```

bash-3.00# /u01/app/11.2.0/grid/bin/crsctl stat res ora.DATA1.dg -p | grep START_DEPENDENCIES
START_DEPENDENCIES=hard(ora.asm,sun.scal-asmdg1-rs) pullup(ora.asm)

bash-3.00# clresource show -p resource_dependencies_offline_restart asm-data1-rs
=== Resources ===

Resource: asm-data1-rs
Resource_dependencies_offline_restart: asm-inst-rs scal-asmdg1-rs

— Standard and extension properties —

```

2. Set the new dependency.

Note that the command includes the minus (-) symbol.

```

bash-3.00# clresource set -p resource_dependencies_offline_restart=-scal-asmdg1-rs asm-data1-rs

```

3. Verify the modified dependency.

```

bash-3.00# /u01/app/11.2.0/grid/bin/crsctl stat res ora.DATA1.dg -p | grep START_DEPENDENCIES
START_DEPENDENCIES=hard(ora.asm) pullup(ora.asm)

bash-3.00# clresource show -p resource_dependencies_offline_restart asm-data1-rs
=== Resources ===

Resource: asm-data1-rs
Resource_dependencies_offline_restart: asm-inst-rs

— Standard and extension properties —

```

Top

How to Delete the `sun.resource` Resource

1. Become superuser.
2. Ensure that the dependency is removed, as described in [How to Remove a dependency](#), and that `sun.resource` is stopped.

```

bash-3.00# /u01/app/11.2.0/grid/bin/crsctl stop res sun.scal-asmdg1-rs
CRS-2500: Cannot stop resource 'sun.scal-asmdg1-rs' as it is not running
CRS-4000: Command Stop failed, or completed with errors.

```

3. Delete `sun.resource`.

```

bash-3.00# /u01/app/11.2.0/grid/bin/crsctl delete res sun.scal-asmdg1-rs

```

4. Verify that `sun.resource` is deleted.

```

bash-3.00# /u01/app/11.2.0/grid/bin/crsctl stat res sun.scal-asmdg1-rs -p
CRS-210: Could not find resource 'sun.scal-asmdg1-rs'.

```

Top

Data Service for Oracle Real Application Clusters Guide

This section discusses errors, omissions, and additions in the Sun Cluster Data Service for Oracle Real Application Clusters Guide.

Configuring an Oracle RAC Instance Proxy Resource Group Without Oracle ASM (CLI)

For Oracle 11g release 2, a patch and additional configuration procedures are required to support using Sun Cluster maintenance commands to configure an Oracle RAC instance proxy resource group without Oracle ASM. Perform the following tasks to add this support:

Task	Additional Information
1. Apply at least the following minimum HA-Oracle patch version on each node that will run Oracle 11g release 2: <ul style="list-style-type: none"> • SPARC: 126047-16 • x86: 126048-16 	
2. Configure the disk set or disk group to use.	<p>Follow one of the following procedures in the Sun Cluster Data Service for Oracle Real Application Clusters Guide:</p> <ul style="list-style-type: none"> • If you are using Solaris Volume Manager for Sun Cluster, create a multi-owner disk set. Follow procedures in How to Create a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster for the Oracle RAC Database . • If you are using Veritas Volume Manager with the cluster feature, create a shared disk group. Follow procedures in How to Create a VxVM Shared-Disk Group for the Oracle RAC Database .
3. Perform procedures in Creating an Oracle Grid Infrastructure Resource .	<p>When you configure Oracle RAC 11g release 2, this procedure replaces How to Create an Oracle CRS Resource for Interoperation With Sun Cluster in Sun Cluster Data Service for Oracle Real Application Clusters Guide.</p> <p>You can alternatively perform this task by using the Sun Cluster Manager GUI or <code>clsetup</code> wizards, as described in How to Enable Sun Cluster and Oracle 10g Release 2 or 11g CRS to Interoperate . If you do, and a patch for CR 6995923 is not yet available, you must first perform Step 1 through Step 3 of this procedure, to create the prerequisite Oracle Grid Infrastructure resource, before you run the wizard. The wizard does not currently create a missing Oracle Grid Infrastructure resource for you.</p>
4. Perform the appropriate CLI procedure in Sun Cluster Data Service for Oracle Real Application Clusters Guide to create Sun Cluster resources for interoperation with Oracle .	<p>When an offline-restart dependency is set, the VALIDATE method of the <code>SUNW.scalable_rac_server_proxy</code> resource transparently modifies the appropriate Grid Infrastructure database resource as follows:</p> <ul style="list-style-type: none"> • <code>ora.dbname.db --> [add hard-start dependency] --> sun.scal-dg1-rs</code> • <code>ora.dbname.db --> [add hard-start dependency] --> sun.scal-dg2-rs</code>

[Top](#)

Creating an Oracle Grid Infrastructure Resource

Perform this procedure to manually create an Oracle Grid Infrastructure resource that proxies the Sun Cluster `SUNW.ScalDeviceGroup` or `SUNW.ScalMountPoint` resource. In this example procedure, the Oracle Grid Infrastructure resource is named `sun.resource`. You configure `sun.resource` to ensure that the corresponding Oracle database is not started until `sun.resource` is online. The `sun.resource` resource comes online only if the corresponding `SUNW.ScalDeviceGroup` or `SUNW.ScalMountPoint` resource is online. And the `SUNW.ScalDeviceGroup` or `SUNW.ScalMountPoint` resource only comes online if the actual volume-manager disk set or disk group or the mount point is online.

[Top](#)

How to Create an Oracle Grid Infrastructure Resource

Perform this procedure on one node of the cluster.

1. Become superuser.

2. Create the Oracle Grid Infrastructure `sun.storage_proxy.type` resource type.

In this example procedure, the `/export/home/oinstall/app/11.2.0/grid` directory is the Oracle Grid Infrastructure home directory.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl \
add type sun.storage_proxy.type \
-basetype cluster_resource \
-attr \
"ATTRIBUTE=ACTION_SCRIPT,TYPE=string", \
"ATTRIBUTE=HOSTING_MEMBERS,TYPE=string", \
"ATTRIBUTE=CARDINALITY,TYPE=string", \
"ATTRIBUTE=PLACEMENT,TYPE=string", \
"ATTRIBUTE=SCRIPT_TIMEOUT,TYPE=int", \
"ATTRIBUTE=RESTART_ATTEMPTS,TYPE=int", \
"ATTRIBUTE=ACL,TYPE=string", \
"ATTRIBUTE=VERSION,TYPE=string"
```

3. Create the Sun Cluster `sun.resource` resource of type `sun.storage_proxy.type`.

The Oracle Grid Infrastructure resource name uses the form `sun.sc-resource`, where `sc-resource` is the name of the `SUNW.ScalDeviceGroup` or `SUNW.ScalMountPoint` resource. For example, the Sun Cluster resource `scal-dg1-rs` is identified as `sun.scal-dg1-rs` in Oracle Grid Infrastructure.

You must also set the following attributes to values that are appropriate for your cluster:

- `HOSTING_MEMBERS` = The list of nodes in the cluster membership
- `CARDINALITY` = The number of nodes in the cluster membership

To create the Oracle Grid Infrastructure `storage_proxy` resource, execute the following command, substituting appropriate values for your cluster:

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl add resource sun.scal-dg1-rs \
-type sun.storage_proxy.type \
-attr "ACTION_SCRIPT=/opt/SUNWscor/dsconfig/bin/scproxy_crs_action" \
CARDINALITY='3' \
SCRIPT_TIMEOUT='20' \
PLACEMENT='restricted' \
RESTART_ATTEMPTS='60' \
HOSTING_MEMBERS='pnsx1 pnsx2 pnsx3' \
VERSION='1' "
```

4. Determine the DBA group of the Oracle Grid Infrastructure installation.

In this example procedure, the `/export/home/oinstall/app/11.2.0/grid` directory is the Oracle Grid Infrastructure home directory.

```
# echo ` /export/home/oinstall/app/11.2.0/grid/bin/osdbagrp `
griddba
```

5. Set the primary group of the Oracle Grid Infrastructure `storage_proxy` resource to the group determined in Step 4.

In this example procedure, the `/export/home/oinstall/app/11.2.0/grid` directory is the Oracle Grid Infrastructure home directory.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl setperm resource sun.scal-dg1-rs -g "griddba"
```

6. Determine the DBA group of the Oracle Database Software installation.

In this example procedure, the `/export/home/oinstall/app/oracle/product/11.2.0/db` directory is the Oracle Database Software home directory.

```
# echo `/export/home/oinstall/app/oracle/product/11.2.0/db/bin/osdbagrp`
dba
```

7. Set the group permissions of the Oracle Grid Infrastructure **storage_proxy** resource to the group determined in Step 6.

Omit this step if the Oracle Grid Infrastructure installation DBA group determined in Step 4 and the Oracle Database Software installation DBA group determined in Step 6 are the same DBA group.

In this example procedure, the `/export/home/oinstall/app/11.2.0/grid` directory is the Oracle Grid Infrastructure home directory.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl setperm resource sun.scal-dg1-rs -u "group:dba:r-x"
```

8. Bring online the Oracle Grid Infrastructure **storage_proxy** resource.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl start resource sun.scal-dg1-rs
CRS-2672: Attempting to start 'sun.scal-dg1-rs' on 'pnsx2'
CRS-2672: Attempting to start 'sun.scal-dg1-rs' on 'pnsx3'
CRS-2672: Attempting to start 'sun.scal-dg1-rs' on 'pnsx1'
CRS-2676: Start of 'sun.scal-dg1-rs' on 'pnsx2' succeeded
CRS-2676: Start of 'sun.scal-dg1-rs' on 'pnsx3' succeeded
CRS-2676: Start of 'sun.scal-dg1-rs' on 'pnsx1' succeeded
```

[Top](#)

Removing an Oracle Grid Infrastructure Resource

This section contains the following procedures to remove an Oracle Grid Infrastructure resource.

- [How to Remove a Dependency](#)
- [How to Delete the `sun.resource` Resource](#)

[Top](#)

How to Remove a Dependency

This procedure shows how to set the offline-restart dependency to remove a dependency.

1. Become superuser.
2. Display the current start dependency that the database has on the Oracle Grid Infrastructure **storage_proxy** resource.

In this example procedure, the `/export/home/oinstall/app/11.2.0/grid` directory is the Oracle Grid Infrastructure home directory, the database name is `testdb`, and the Oracle Grid Infrastructure **storage_proxy** resource name is `sun.scal-dg1-rs`.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl stat res ora.testdb.db -p | grep START_DEPENDENCIES
START_DEPENDENCIES=hard(sun.scal-dg1-rs)
weak(type:ora.listener.type,global:type:ora.scan_listener.type,uniform:ora.ons,uniform:ora.eons)

# clresource show -p resource_dependencies_offline_restart rac-server-proxy-rs
=== Resources ===

Resource: rac-server-proxy-rs
Resource_dependencies_offline_restart: crs-fw-rs scal-dg1-rs

— Standard and extension properties —
```

3. Remove the offline-restart dependency on the **SUNW.ScalDeviceGroup** or **SUNW.ScalMountPoint** resource from the Oracle RAC instance proxy resource.

This command clears the dependencies that the Oracle Grid Infrastructure database resource has on the Oracle Grid Infrastructure `storage_proxy` resource. Note that the command includes the minus (-) symbol.

```
# clresource set -p resource_dependencies_offline_restart=-scal-dg1-rs rac-server-proxy-rs
```

4. Verify that the start dependency on the Oracle Grid Infrastructure resource is removed.

In this example procedure, the `/export/home/oinstall/app/11.2.0/grid` directory is the Oracle Grid Infrastructure home directory.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl stat res ora.testdb.db -p | grep START_DEPENDENCIES
START_DEPENDENCIES=weak(type:ora.listener.type,global:type:ora.scan_listener.type,uniform:ora.ons,uniform:ora.eons)

# clresource show -p resource_dependencies_offline_restart rac-server-proxy-rs

=== Resources ===

Resource: rac-server-proxy-rs
Resource_dependencies_offline_restart: crs-fw-rs

— Standard and extension properties —
```

[Top](#)

How to Delete the `sun.resource` Resource

1. Become superuser.
2. Ensure that the dependency is removed, as described in [How to Remove a Dependency](#), and that `sun.resource` is stopped.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl stop res sun.scal-dg1-rs
CRS-2673: Attempting to stop 'sun.scal-dg1-rs' on 'pnsx3'
CRS-2673: Attempting to stop 'sun.scal-dg1-rs' on 'pnsx1'
CRS-2673: Attempting to stop 'sun.scal-dg1-rs' on 'pnsx2'
CRS-2677: Stop of 'sun.scal-dg1-rs' on 'pnsx3' succeeded
CRS-2677: Stop of 'sun.scal-dg1-rs' on 'pnsx1' succeeded
CRS-2677: Stop of 'sun.scal-dg1-rs' on 'pnsx2' succeeded
```

3. Delete `sun.resource`.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl delete res sun.scal-dg1-rs
```

4. Verify that `sun.resource` is deleted.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl stat res sun.scal-dg1-rs -p
CRS-210: Could not find resource 'sun.scal-dg1-rs'.
```

[Top](#)

Man Pages

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

Extraneous Information in Man Pages and Usage Messages

Some man pages and usage messages refer to the use of QFS shared file systems with data services other than Oracle RAC. Support for QFS shared file systems when used with data services other than Oracle RAC has not been enabled in the 3.2 11/09 release. Ignore the information.

[Top](#)

SUNW.Proxy_SMF_scalable(5)

The `SUNW.Proxy_SMF_loadbalanced(5)` man page has been renamed `SUNW.Proxy_SMF_scalable(5)` to reflect the correct resource type name.

[Top](#)

scdpmd.conf(4)

A new man page describes the `scdpmd.conf` file. Editing this file enables you to tune the behavior of the disk-path-monitoring daemon `scdpmd(1M)`. See the `scdpmd.conf(4)` man page.

[Top](#)

clresourcetype(1CL)

The description of the `set` subcommand currently states, "The only resource type property that you can set is the `Installed_nodes` property." This property is no longer the only property that is tunable at any time. See the `rt_properties(5)` man page for the descriptions of resource-type properties, including when each property is tunable.

[Top](#)

r_properties(5)

The following property descriptions have been added to the `r_properties(5)` man page:

CheckActivePortInstances

Boolean that determines whether a node participates in the scalable service, and receives clients requests from the load balancer, when not all ports that are specified in the `Port_list` property have an active listening process.

This property is only applicable to scalable services. Supported values are as follows:

FALSE (default)

A node participates in the scalable service when at least one of its ports has an active listening process.

TRUE

A node participates in the scalable service only when all ports have an active listening process.

- Conditional/Optional
- Default: `FALSE`
- Tunable: When disabled

Outgoing_Connection

Boolean that specifies whether the scalable service uses the virtual network address (see `Network_resources_used` property) in initiating outgoing requests to servers outside the cluster. The load balancer ensures that any incoming replies are forwarded to the initiating node.



Note

Only one node of the cluster can initiate requests to a given server at a time.

This property is only applicable to scalable services with `Generic_Affinity` set to `TRUE` and `Load_balancing_policy` set

to `LB_STICKY_WILD`. Supported values are as follows:

`FALSE` (default)

The scalable service does not initiate outgoing requests to external servers by using the virtual network address that is specified in the `Network_resources_used` property.

`TRUE`

The scalable service initiates outgoing requests to external servers by using the virtual network address that is specified in the `Network_resources_used` property. The load balancer forwards incoming replies to the initiating node.

- Conditional/Optional
- Default: `FALSE`
- Tunable: At Creation

IPsec_Session_Failover

Boolean that specifies whether IP Security associations that pertain to the scalable service can be replicated over the interconnect to other cluster nodes for faster failover.

- IPsec associations that are actively used are not replicated due to the extremely high load it would impose on the cluster interconnect. Only associations that are idle for over 2 minutes by default are replicated.
- To ensure protection of the IPsec state, including keying material during replication, you must deploy IPsec on the interconnect.
- You must use Internet Key Exchange (IKE) for the scalable service to support the replication of IPsec associations.

This property is only applicable to scalable services with `Generic_Affinity` set to `TRUE` and `Load_balancing_policy` set to `LB_STICKY_WILD`. Supported values are:

`FALSE` (default)

No IPsec association state is replicated.

`TRUE`

Perform replication of IPsec association state pertaining to the scalable service over the interconnect to all other cluster nodes. Any node taking over client requests for a failed node immediately makes available such associations for the clients.

- Conditional/Optional
- Default: `FALSE`
- Tunable: At Creation

[Top](#)