

Sun Cluster 3.2 2/08 Release Notes

3.2 2/08 Release Notes

820-2563

Release Notes

Copyright © 2007, 2008, 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 2-08 Release Notes

Sun Cluster 3.2 2/08 Release Notes for Solaris OS

This document provides the following information for Sun™ Cluster 3.2 2/08 software.

- [What's New in the Sun Cluster 3.2 2/08 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 2/08 Documentation](#)
- [Documentation Issues](#)
- [Resolved Issues](#)

[Top](#)

What's New in the Sun Cluster 3.2 2/08 Software

This section provides information related to new features, functionality, and supported products in the Sun Cluster 3.2 2/08 software. This section also provides information on any restrictions that are introduced in this release.

[Top](#)

New Features and Functionality

This section describes new features and functionality in the Sun Cluster 3.2 2/08 software.

The following new features and functionality are provided in patches to Sun Cluster 3.2 2/08 software.

- [Support for Exclusive-IP Zones](#)
- [New `Global_zone_override` resource property](#)
- [\(SPARC\) Support for solaris9 Brand Zones in HA-Containers](#)
- [\(SPARC\) Support for Oracle 11g](#)
- [\(SPARC\) Support for Logical Domains \(LDom\) Guest Domains as Cluster Nodes](#)

The following new features and functionality are provided in the initial Sun Cluster 3.2 2/08 release.

- [Sun Service Tags Support](#)
- [Support for Using NAS Devices From Sun Microsystems as Shared Storage and Quorum Devices](#)
- [Support for EMC Symmetrix Remote Data Facility \(SRDF\)](#)
- [HA-Containers Support for Solaris Zones Brands](#)
- [Editable Properties in Sun Cluster Manager](#)
- [New Sun Cluster Upgrade Guide](#)
- [Quorum Server Documentation Moved to the Installation and System Administration Guides](#)

[Top](#)

Support for Exclusive-IP Zones

Sun Cluster software now supports non-global zones of `ip-type=exclusive`. This functionality is introduced in patch 126106-17 for SPARC based platforms and 126107-17 for x86 based platforms. To run logical-hostname resources in exclusive-IP zones, you must register version 3 of the `SUNW.LogicalHostname` resource type after you apply this patch.

1. Determine whether version 3 of the `SUNW.LogicalHostname` resource type is already registered.

```
phys-schost# clresource type list SUNW.LogicalHostname:3
```

2. If `SUNW.LogicalHostname:3` is not registered, run the following command.

```
phys-schost# clresource type register SUNW.LogicalHostname
```

For more details, see ["Registering a Resource Type"](#) in Sun Cluster Data Services Planning and Administration Guide.

The following information provides guidelines and procedures for using exclusive-IP zones in a Sun Cluster configuration.

[Top](#)

Guidelines for Exclusive-IP Zones

- Logical-hostname resource groups - In a resource group that contains a `LogicalHostname` resource, if the node list contains any non-global zone with the `ip-type` property set to `exclusive`, then all zones in that node list must have this property set to `exclusive`. Note that a global zone always has the `ip-type` property set to `shared`, and therefore cannot coexist in a node list that contains zones of `ip-type=exclusive`. This restriction applies only to versions of the Solaris OS that use the Solaris zones `ip-type` property.
- IPMP groups - For all public-network adapters that are used for data-service traffic in the non-global zone, you must manually configure IPMP groups in all `/etc/hostname.adapter` files on the zone. This information is not inherited from the global zone. For guidelines and instructions to configure IPMP groups, follow the procedures in [Part VI, IPMP](#), in [System Administration Guide: IP Services](#).
- Private-hostname dependency - Exclusive-IP zones cannot depend on cluster private hostnames and addresses.
- Shared-address resources - Shared-address resources cannot use exclusive-IP zones.

[Top](#)

Configuring IPMP Groups for Exclusive-IP Zones

After you create an exclusive-IP non-global zone in a Sun Cluster configuration, perform the following steps.

1. Manually configure IPMP groups in each `/etc/hostname.adapter` file that is on the zone. You must configure an IPMP group for each public-network adapter that is used for data-service traffic in the zone. This information is not inherited from the global zone. See ["Public Networks"](#) in the Sun Cluster Software Installation Guide for more information about configuring IPMP groups in a cluster.
2. Set up name-to-address mappings for all logical-hostname resources that are used by the zone.
 - a. Add name-to-address mappings to the `/etc/inet/hosts` file on the zone. This information is not inherited from the global zone.
 - b. If you use a name server, add the name-to-address mappings.

[Top](#)

New `Global_zone_override` Resource Property

A new resource property, `Global_zone_override`, is introduced in patch 126106-17 for SPARC based platforms and 126107-17 for x86 based platforms. This property overrides the value of the resource type property `Global_zone` for a specified resource. For more information about the `Global_zone_override` resource property and its effect on the `Global_zone` resource type property, see the sections for the [r_properties\(5\)](#) and [rt_properties\(5\)](#) man pages.

[Top](#)

(SPARC) Support for solaris9 Brand Zones in HA-Containers

The Sun Cluster Data Service for Solaris Containers now supports `solaris9` branded zones on SPARC based systems. The minimum patch requirement for this functionality is 126020-03.

Procedures in the [Sun Cluster Data Service for Solaris Containers Guide](#) that pertain to `solaris8` branded zones also apply to `solaris9` branded zones.

[Top](#)

(SPARC) Support for Oracle 11g

Sun Cluster software now supports Oracle 11g in HA-Oracle and Oracle RAC configurations on SPARC based platforms. Support is introduced in the May 2008 Core patch. Procedures in the Sun Cluster 3.2 2/08 documentation set are valid for Oracle 11g. If an instruction is specific to a particular Oracle version, use instructions that apply to Oracle 10g for an Oracle 11g configuration.

**Note:**

When you install the `ORCLudlm` package on Sun Cluster nodes for Oracle 11g RAC support, ensure that the package is installed with a UNIX group that is part of the groups of the `crs` user (Oracle Clusterware software owner) and the groups of the `oracle` user (Oracle Database software owner). For example, if the `dba` group is used for `ORCLudlm` installation, the `dba` group should be either the primary group or the secondary group for both the `crs` user and the `oracle` user.

[Top](#)

(SPARC) Support for Logical Domains (LDoms) Guest Domains as Cluster Nodes

You can now configure Logical Domains (LDoms) 1.0.3 guest domains as virtual Sun Cluster nodes. In this configuration, a guest domain node is viewed the same as a physical node of a cluster. Support is introduced in the May 2008 Core patch.

You can configure Logical Domains (LDoms) guest and I/O domains as virtual Sun Cluster nodes. In other words, you can create a clustered pair, pair+N, N+1, and N*N cluster that consists of any combination of physical machines, LDoms I/O domains, and LDoms guest domains. You can also create clusters that consist of only LDoms I/O domains or only guest domains.

The following sections describe topologies, requirements and restrictions, and installation and configuration procedures for guest domains as cluster nodes:

- [SPARC: LDom Guest Domains: Cluster in a Box Topology](#)
- [SPARC: LDoms Guest Domains: Single Cluster Spans Two Different Servers Topology](#)
- [SPARC: LDoms Guest Domains: Clusters Span Two Different Servers Topology](#)
- [SPARC: LDom Guest Domains: Redundant I/O Domains](#)
- [SPARC: Software and Patch Requirements to Support LDoms Software](#)
- [SPARC: Guidelines for Logical Domains in a Cluster](#)
- [SPARC: How to Install Logical Domains Software and Create Domains](#)
- [SPARC: Configuring Logical Domains as Cluster Nodes](#)

[Top](#)

SPARC: LDom Guest Domains: Cluster in a Box Topology

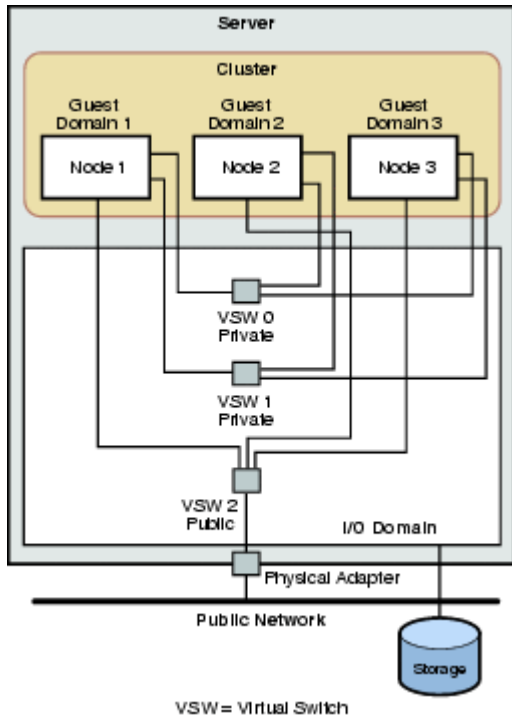
In this Logical Domains (LDoms) guest domain topology, a cluster and every node within that cluster are located on the same server. Each LDoms guest domain node acts the same as a physical node in a cluster. To preclude your having to include a quorum device, this configuration includes three nodes rather than only two.

In this topology, you do not need to connect each virtual switch (vsw) for the private network to a physical network because they need only communicate with each other. In this topology, cluster nodes can also share the same storage device, as all cluster nodes are located on the same server. To learn more about guidelines for using and installing LDoms domains in a cluster, see [SPARC: Guidelines for Logical Domains in a Cluster](#), [SPARC: How to Install Logical Domains Software and Create Domains](#), and [SPARC: Configuring LDoms Domains as Cluster Nodes](#).

This topology does not provide high availability, as all nodes in the cluster are located on the same server. However, developers and administrators might find this topology useful for testing and other non-production tasks. This topology is also called a "cluster in a box".

The following figure illustrates a cluster in a box configuration.

SPARC: Cluster in a Box Topology



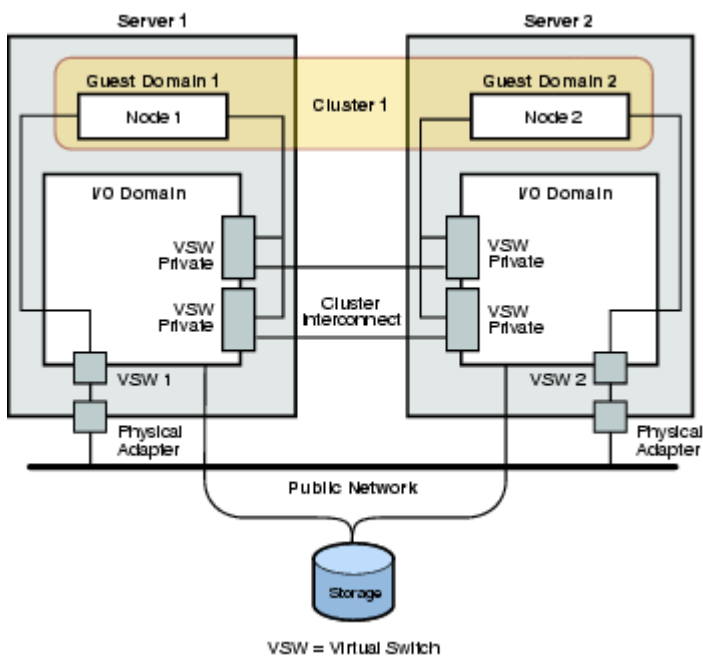
Top

SPARC: LDom's Guest Domains: Single Cluster Spans Two Different Servers Topology

In this Logical Domains (LDoms) guest domain topology, a single cluster spans two different servers and each cluster comprises one node on each server. Each LDom's guest domain node acts the same as a physical node in a cluster. To learn more about guidelines for using and installing LDoms domains in a cluster, see [SPARC: Guidelines for Logical Domains in a Cluster](#), [SPARC: How to Install Logical Domains Software and Create Domains](#), and [SPARC: Configuring LDoms Domains as Cluster Nodes](#).

The following figure illustrates a configuration in which a single cluster spans two different servers.

SPARC: Single Cluster Spans Two Different Servers



Top

SPARC: LDom's Guest Domains: Clusters Span Two Different Servers Topology

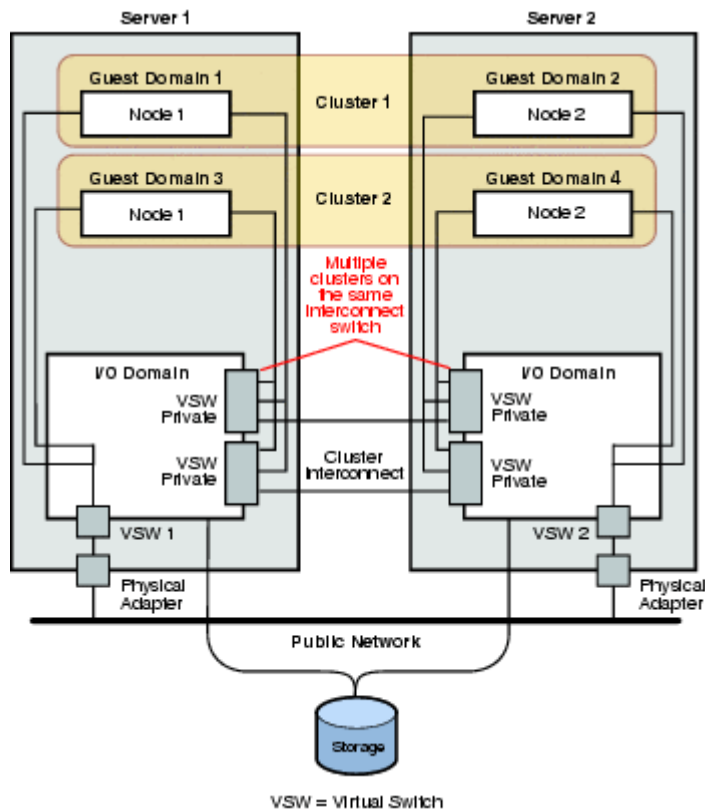
In this Logical Domains (LDoms) guest domain topology, each cluster spans two different servers and each cluster comprises one

node on each server. Each LDom's guest domain node acts the same as a physical node in a cluster. In this configuration, because both clusters share the same interconnect switch, you must specify a different private network address on each cluster. Otherwise, if you specify the same private network address on clusters that share an interconnect switch, the configuration fails.

To learn more about guidelines for using and installing LDom's domains in a cluster, see [SPARC: Guidelines for Logical Domains in a Cluster](#), [SPARC: How to Install Logical Domains Software and Create Domains](#), and [SPARC: Configuring LDom's Domains as Cluster Nodes](#).

The following figure illustrates a configuration in which more than a single cluster spans two different servers.

SPARC: Clusters Span Two Different Servers



[Top](#)

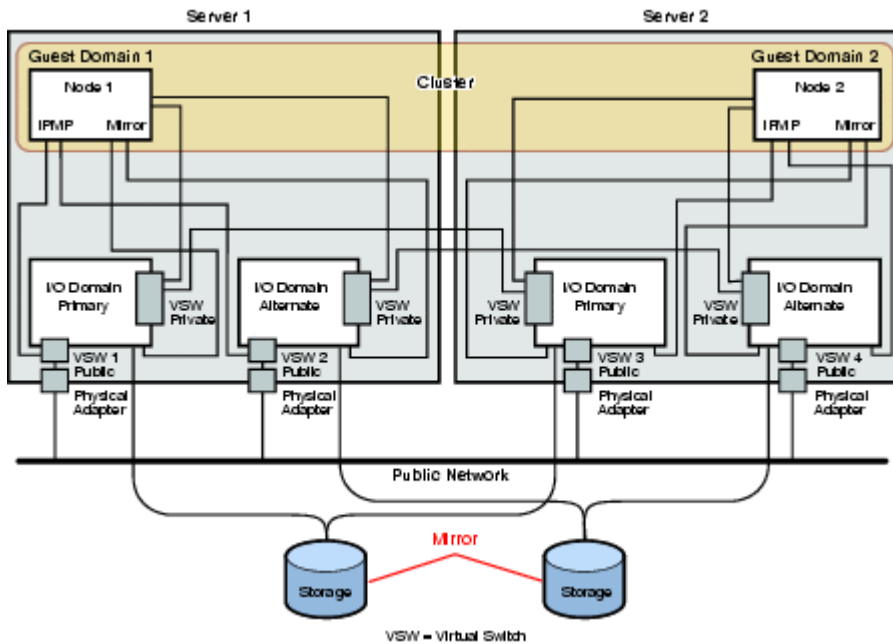
SPARC: LDom Guest Domains: Redundant I/O Domains

In this Logical Domains (LDoms) guest domain topology, multiple I/O domains ensure that guest domains, or nodes within the cluster, continue to operate if an I/O domain fails. Each LDom's guest domain node acts the same as a physical node in a cluster.

In this topology, the guest domain runs IP network multipathing (IPMP) across two public networks, one through each I/O domain. Guest domains also mirror storage devices across different I/O domains. To learn more about guidelines for using and installing LDom's domains in a cluster, see [SPARC: Guidelines for Logical Domains in a Cluster](#), [SPARC: How to Install Logical Domains Software and Create Domains](#), and [SPARC: Configuring LDom's Domains as Cluster Nodes](#).

The following figure illustrates a configuration in which redundant I/O domains ensure that nodes within the cluster continue to operate if an I/O domain fails.

SPARC: Redundant I/O Domains



[Top](#)

SPARC: Software and Patch Requirements to Support LDom Software

Sun Cluster 3.2 2/08 software supports LDom 1.0.3 I/O domains and guest domains as cluster nodes, with the following requirements:

- Solaris 10 5/08 OS
- At least Sun Cluster patch version 126106-15, installed in the guest domain
- The following Solaris 10 patches (minimum versions):
 - 137042-01
 - 137111-01
 - 138042-02
 - 138056-01

The Solaris 10 patches must be installed in the following locations:

- Guest domains that run Sun Cluster software
- Any I/O domains that export virtual devices to guest domains that run Sun Cluster software

[Top](#)

SPARC: Guidelines for Logical Domains in a Cluster

The following is new information to support Logical Domains (LDoms) I/O domains or guest domains on a physically clustered machine that is SPARC hypervisor capable:

- SCSI LUN requirement - The virtual shared storage device, or virtual disk backend, of an LDom guest domain must be a full SCSI LUN in the I/O domain. You cannot use an arbitrary virtual device.
- Fencing - Do not export a storage LUN to more than one guest domain on the same physical machine, unless you also disable fencing for that device. Otherwise, if two different guest domains on the same machine both have visibility to a device, the device would be fenced whenever one of the guest domains dies. The fencing of the device would panic any other guest domain that subsequently tries to access the device.
- Network isolation- Guest domains that are located on the same physical machine but are configured in different clusters must be network isolated from each other. Use one of the following methods:
 - Configure the clusters to use different network interfaces in the I/O domain for the private network.
 - Use different network addresses for each of the clusters.
- Networking in guest domains- Network packets to and from guest domains must traverse service domains to reach the network drivers through virtual switches. Virtual switches use kernel threads that run at system priority. The virtual switch threads must be able to acquire needed CPU resources to perform critical cluster operations, including heartbeats, membership, checkpoints, and so forth.

While configuring virtual switches with the `mode=sc` setting enables expedited handling of cluster heartbeat packets, the reliability of other critical cluster operations can be enhanced by adding more CPU resources to the service domain under the following workloads:

- High interrupt load, such as due to network or disk I/O. Under extreme load, virtual switches can preclude system threads from running for a long time, including virtual switch threads.
- Real-time threads that are overly aggressive in retaining CPU resources. Real-time threads run at a higher priority than virtual-switch threads, which can restrict CPU resources for virtual-switch threads for an extended time.
- Exporting storage from I/O domains- If you configure a cluster that is composed of Logical Domains I/O domains, do not export its storage devices to other guest domains that also run Sun Cluster software.
- Solaris I/O multipathing - Do not run Solaris I/O multipathing software (MPxIO) from guest domains. Instead, run Solaris I/O multipathing software in the I/O domain and export it to the guest domains.
- Private interconnect IP address range - The private network is shared by all guest domains that are created on the same physical machine and is visible to all these domains. Before you specify a private-network IP address range to the `scinstall` utility for use by a guest-domain cluster, ensure that the address range is not already in use by another guest domain on the same physical machine.
- Adapter names - When you configure an I/O domain or a guest domain as a cluster node, specify adapter names by their virtual names, `vnetN`, such as `vnet0`, `vnet1`, and so forth. Virtual adapter names are recorded in the `/etc/path_to_inst` file.
- Cluster Control Panel - You cannot use Cluster Control Panel (CCP) software, such as `cconsole` or `cssh`, to connect to LDoms guest domains. However, you can use CCP software to connect to LDoms I/O domains.

For more information about Logical Domains, see the [Logical Domains \(LDoms\) 1.0.3 Administration Guide](#).

[Top](#)

SPARC: How to Install Logical Domains Software and Create Logical Domains

Perform this procedure to install Logical Domains (LDoms) software on a physically clustered machine and to create I/O domains and guest domains.

Before You Begin

Perform the following tasks:

- Ensure that the machine is SPARC hypervisor capable.
- Ensure that the Solaris 10 OS is installed.
- Have available [Logical Domains \(LDoms\) 1.0.3 Administration Guide](#) and [Logical Domains \(LDoms\) 1.0.3 Release Notes](#).
- Read the requirements and guidelines in [SPARC: Guidelines for Logical Domains in a Cluster](#).

Procedures

1. Become superuser on the machine.
2. Install Logical Domains software and configure logical domains.
 - Follow procedures in [Installing and Enabling Software](#) in [Logical Domains \(LDoms\) 1.0.3 Administration Guide](#). If you create guest domains, adhere to Sun Cluster guidelines for creating guest domains in a cluster.
 - Use the `mode=sc` option for all virtual switch devices that connect the virtual network devices that are used as the cluster interconnect.
 - For shared storage, map only the full SCSI disks into the guest domains.
 - For virtual disks that you add to a guest domain, set the timeout value to 600 seconds. If you later change this value for tuning purposes, set the timeout to a minimum of 30 seconds to avoid communication failures between the I/O domain and the guest domain.
3. Follow procedures in [How to Install Sun Cluster Framework and Data-Service Software Packages](#) in [Sun Cluster Software Installation Guide](#) to install Sun Cluster software in the LDoms I/O domains or guest domains.

[Top](#)

SPARC: Configuring Logical Domains as Cluster Nodes

To configure one or more LDoms guest domains or I/O domains as cluster nodes, perform procedures in [Establishing the Cluster](#) in [Sun Cluster Software Installation Guide](#), as you would to configure physical machines as cluster nodes. Observe the following special instructions for logical domains:

- The nodes of a cluster can consist of any combination of physical machines, LDoms I/O domains, and LDoms guest

domains.

- If a physical machine is configured with logical domains, install Sun Cluster software only in I/O domains or guest domains on that machine.
- Specify adapters by the virtual name, `vnet.N`.

[Top](#)

Sun Service Tags Support

This release supports Sun Service Tags. Service tags provide the customer with an easy mechanism for identifying their products to Sun. For more information about Sun Service Tags, see Sun Connection on BigAdmin (<http://www.sun.com/bigadmin/hubs/connection/tasks/register.jsp>).

[Top](#)

Support for Using NAS Devices From Sun Microsystems as Shared Storage and Quorum Devices

You can now use the following network-attached storage (NAS) devices from Sun Microsystems, Inc. as shared storage and quorum devices in a Sun Cluster environment.

- Sun StorEdge™ NAS Appliance
- Sun StorageTek™ NAS Appliance



Note

A Sun NAS device is supported as a quorum device only in a two-node cluster.

[Top](#)

Support for EMC Symmetrix Remote Data Facility (SRDF)

Sun Cluster software now supports the use of EMC Symmetrix Remote Data Facility Software (SRDF) for storage-based data replication, as used in a campus cluster. The addition of this new support provides the user with two choices of storage-based data replication solutions for use with Sun Cluster:

- Hitachi True Copy
- EMC SRDF

[Top](#)

HA-Containers Support for Solaris™ Zones Brands

In this release, the Sun Cluster Data Service for Solaris Containers supports the following branded zones:

- SPARC: `native`, `solaris8`
- x86: `native`, `lx`

[Top](#)

Editable Properties in Sun Cluster Manager

Enhancements have been made to increase the usability of editable properties in Sun Cluster Manager.

[Top](#)

New Sun Cluster Upgrade Guide

Upgrade guidelines and procedures have been moved from the Sun Cluster Software Installation Guide for Solaris OS to a new manual, [Sun Cluster Upgrade Guide for Solaris OS](#).

[Top](#)

Quorum Server Documentation Moved to the Installation and System Administration Guides

Information in the Sun Cluster Quorum Server User's Guide for the Sun Cluster 3.2 release has been moved into other manuals for the Sun Cluster 3.2 2/08 release.

- Information to install a quorum server is now located in "[How to Install and Configure Quorum Server Software](#)" in Sun Cluster Software Installation Guide for Solaris OS.
- Information to administer a quorum server is now located in "[Administering Sun Cluster Quorum Servers](#)" in Sun Cluster System Administration Guide for Solaris OS.

[Top](#)

Features Nearing End of Life

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

No features are nearing end of life at this time.

[Top](#)

Compatibility Issues

- [SPARC: Logical Domains \(LDoms\)](#)
- [Exclusive-IP Zones](#)
- [fssnap Support](#)
- [Solaris Volume Manager GUI](#)
- [Loopback File System \(LOFS\)](#)
- [Shared QFS With Solaris Volume Manager for Sun Cluster](#)
- [Solaris Trusted Extensions](#)
- [Accessibility Features for People With Disabilities](#)

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](#)

SPARC: Logical Domains (LDoms)

CR 6684721 (Need a way to force file backed virtual I/O to be synchronous) is a known issue in Logical Domains (LDoms) 1.0.3 software that impacts Sun Cluster software. The workaround is to add the following entry to the `/etc/system` file on the I/O domain:

```
set vds:vd_file_write_flags = 0
```

Contact your Sun service representative to determine whether a patch is available.

Exclusive-IP Zones

The use of LogicalHostname and SharedAddress resources is limited to non-global zones of shared-IP type only. This is the default IP type for zones and no administrator steps are required.

You can still use zones of `ip-type=exclusive` with the cluster, but network resources, including NICs, IPMP groups, and IP addresses, are not monitored and managed by Sun Cluster software.

[Top](#)

`fsnap` Support

Sun Cluster software does not support the `fsnap` command, which is a feature of UFS. However, you can use the `fsnap` command on local file systems that are not controlled by Sun Cluster software. The following restrictions apply to `fsnap` support:

- The `fsnap` command is supported on local file systems that are not managed by Sun Cluster software.
- The `fsnap` command is not supported on cluster file systems.
- The `fsnap` command is not supported on local file systems that are under the control of HAStoragePlus.

[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](#)

Loopback File System (LOFS)

Sun Cluster 3.2 2/08 software does not support the use of LOFS under certain conditions. If you must enable LOFS on a cluster node, such as when you configure non-global zones, first determine whether the LOFS restrictions apply to your configuration. See the guidelines in "[Solaris OS Feature Restrictions](#)" in Sun Cluster Software Installation Guide for Solaris OS for more information about the restrictions and workarounds that permit the use of LOFS when restricting conditions exist.

[Top](#)

Shared QFS With Solaris Volume Manager for Sun Cluster

The following behavior has been reported on x86 Oracle RAC configurations with one or more Sun StorageTek QFS shared file systems that mount devices from multiowner disksets of Solaris Volume Manager for Sun Cluster (CR 6655081).

If the QFS file system metadata server is located on a node that is not the master node of the disk set and that node loses all storage connectivity, Oracle CRS will reboot the node. At the same time, the other nodes that are QFS metadata client nodes might experience errors writing to related database files. The write error condition corrects itself when the Oracle RAC instances are restarted. This restart should be an automatic recovery action by Oracle CRS, after the QFS metadata server and Solaris Volume Manager automatic recovery has been completed by Sun Cluster software.

[Top](#)

Solaris Trusted Extensions

Sun Cluster 3.2 2/08 software supports Solaris non-global zones within a cluster. Starting with the Solaris 10 11/06 release, the Solaris 10 OS includes support for Solaris Trusted Extensions. Solaris Trusted Extensions uses non-global zones as well. The interaction between Sun Cluster and Solaris Trusted Extensions by using non-global zones has not been tested. Customers are advised to proceed with caution when using these technologies.

[Toph3. 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](#)

Commands Modified in This Release

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

[Top](#)

Object-Oriented Command-Line Interface

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

[Top](#)

Resource Properties

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

[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 2/08 software is distributed under Solaris Cluster 3.2 2/08 and Sun Java Availability Suite.

|| Current Product Name || Former Product Name ||

Sun Cluster Manager	SunPlex Manager
Sun Cluster Agent Builder	SunPlex Agent Builder
Sun Cluster HA for MaxDB	Sun Cluster HA for SAP DB

[Top](#)

Supported Products

This section describes the supported software and memory requirements for Sun Cluster 3.2 2/08 software.

- Memory requirements - Sun Cluster 3.2 2/08 software requires the following memory requirements for every cluster node:
 - Minimum of 512 Mbytes 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.

- RSM-API - Sun Cluster 3.2 2/08 software supports the Remote Shared Memory Application Programming Interface (RSM-API) on RSM-capable interconnects, such as PCI-SCI.
- Solaris Operating System (OS) - Sun Cluster 3.2 2/08 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 11/06, Solaris 10 8/07, Solaris 10 5/08*, Solaris 10 10/08*
 - *The Solaris 10 5/08 and Solaris 10 10/08 OS require the latest Sun Cluster patches.



Note

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

- Volume managers

Platform	Operating System	Volume Manager	Cluster Feature
SPARC	Solaris 9	Solaris Volume Manager	Solaris Volume Manager for Sun Cluster
"	"	Veritas Volume Manager components that are delivered as part of Veritas Storage Foundation 5.0. This support requires VxVM 5.0 MP1.	Veritas Volume Manager 5.0 cluster feature (with RAC only)
SPARC	Solaris 10	Solaris Volume Manager	Solaris Volume Manager for Sun Cluster
"	"	Veritas Volume Manager components that are delivered as part of Veritas Storage Foundation 5.0. This support requires VxVM 5.0 MP1.	Veritas Volume Manager 5.0 cluster feature (with RAC only)
x86	Solaris 10	Solaris Volume Manager	Solaris Volume Manager for Sun Cluster
"	"	Veritas Volume Manager components that are delivered as part of Veritas Storage Foundation 5.0	Not applicable - Sun Cluster 3.2 2/08 software does not support the VxVM cluster feature on the x86 platform.

- File systems

Platform	Operating System	File System	Features and External Volume Management
SPARC	Solaris 9	Solaris UFS	Not applicable
"	"	Sun StorEdge QFS	Not applicable
"	"	QFS 4.5 - Standalone file system	<ul style="list-style-type: none"> Features: HA-NFS, HA-Oracle External Volume Management: Solaris Volume Manager, VxVM
"	"	QFS 4.5 - Shared QFS file system	<ul style="list-style-type: none"> Feature: Oracle RAC External Volume Management: Solaris Volume Manager for Sun Cluster
"	"	QFS 4.6	<ul style="list-style-type: none"> Features: COTC-Shared QFS clients outside the cluster, HA-SAM Failover
"	"	Veritas File System components that are delivered as part of Veritas Storage Foundation 5.0	Not applicable
SPARC	Solaris 10	Solaris UFS	Not applicable
"	"	Solaris ZFS	Not applicable
"	"	Sun StorEdge QFS	Not applicable
"	"	QFS 4.5 - Standalone file system	<ul style="list-style-type: none"> Features: HA-NFS, HA-Oracle External Volume Management: Solaris Volume Manager, VxVM
"	"	QFS 4.5 - Shared QFS file system	<ul style="list-style-type: none"> Feature: Oracle RAC External Volume Management: Solaris Volume Manager for Sun Cluster

"	"	QFS 4.6	<ul style="list-style-type: none"> • Features: COTC-Shared QFS clients outside the cluster, HA-SAM Failover
"	"	Veritas File System components that are delivered as part of Veritas Storage Foundation 5.0	Not applicable
x86	Solaris 10	Solaris UFS	Not applicable
"	"	Solaris ZFS	Not applicable
"	"	Sun StorEdge QFS	Not applicable
"	"	QFS 4.5 - Standalone file system	<ul style="list-style-type: none"> • Features: HA-NFS, HA-Oracle • External Volume Management: Solaris Volume Manager, VxVM
"	"	QFS 4.5 - Shared QFS file system	<ul style="list-style-type: none"> • Feature: Oracle RAC • External Volume Management: Solaris Volume Manager for Sun Cluster
"	"	QFS 4.6	<ul style="list-style-type: none"> • Features: COTC-Shared QFS clients outside the cluster, HA-SAM Failover
"	"	Veritas File System components that are delivered as part of Veritas Storage Foundation 5.0	Not applicable

- Sun StorEdge™ Availability Suite 3.2.1 (Solaris 9 only) and Sun StorageTek™ Availability Suite 4.0 (Solaris 10 only)
- Sun Management Center 3.6.1 and 4.0
- Logical Domains (LDoms) 1.0.3 - Sun Cluster 3.2 2/08 software supports LDoms 1.0.3 I/O domains and guest domains as cluster nodes. Also see [SPARC: Software and Patch Requirements to Support LDoms Software](#).
- Common agent container 2.1
- Data services (agents) - Contact your Sun sales representative for the complete list of supported data services and application versions.



Note

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

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

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

- Sun Cluster Data Service for Sun Java System Web Server
- Sun Cluster Data Service for SWIFTAlliance Access
- Sun Cluster Data Service for SWIFTAlliance Gateway
- Sun Cluster Data Service for Sybase ASE



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.

The following data service is not supported on Solaris 10 OS in this Sun Cluster release.

- Sun Cluster Data Service for Agfa IMPAX

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

Data Service	Sun Cluster Resource Type
Sun Cluster HA for Agfa IMPAX	SUNW.gds
Sun Cluster HA for Apache	SUNW.apache
Sun Cluster HA for Apache Tomcat	SUNW.gds
Sun Cluster HA for DHCP	SUNW.gds
Sun Cluster HA for DNS	SUNW.dns
Sun Cluster HA for MaxDB	SUNW.sapdb, SUNW.sap_xserver
Sun Cluster HA for MySQL	SUNW.gds
Sun Cluster HA for NFS	SUNW.nfs
Sun Cluster Oracle Application Server	SUNW.gds
Sun Cluster HA for Oracle E-Business Suite	SUNW.gds
Sun Cluster HA for Oracle	SUNW.oracle_server, SUNW.oracle_listener
Sun Cluster Support for Oracle Real Application Clusters	SUNW.rac_framework, SUNW.rac_udlm, SUNW.rac_svm, SUNW.rac_cvm, SUNW.oracle_rac_server, SUNW.oracle_listener, SUNW.scaldevicegroup, SUNW.scalmountpoint, SUNW.crs_framework, SUNW.scalable_rac_server_proxy
Sun Cluster HA for PostgreSQL	SUNW.gds
Sun Cluster HA for Samba	SUNW.gds
Sun Cluster HA for SAP	SUNW.sap_ci, SUNW.sap_ci_v2, SUNW.sap_as, SUNW.sap_as_v2
Sun Cluster HA for SAP liveCache	SUNW.sap_livecache, SUNW.sap_xserver
Sun Cluster HA for SAP Web Application Server	SUNW.sapeng, SUNW.saprepl, SUNW.sapscs, SUNW.sapwebas
Sun Cluster HA for Siebel	SUNW.sblgtwy, SUNW.sblsrvr
Sun Cluster HA for Solaris Containers	SUNW.gds
Sun Cluster HA for N1 Grid Engine	SUNW.gds

Sun Cluster HA for Sun Java System Application Server supported versions before 8.1	SUNW.slas
Sun Cluster HA for Sun Java System Application Server supported versions as of 8.1	SUNW.jsas, SUNW.jsas-na
Sun Cluster HA for Sun Java System Application Server EE (supporting HADB versions before 4.4)	SUNW.hadb
Sun Cluster HA for Sun Java System Application Server EE (supporting HADB versions as of 4.4)	SUNW.hadb_ma
Sun Cluster HA for Sun Java System Message Queue	SUNW.s1mq
Sun Cluster HA for Sun Java System Web Server	SUNW.iws
Sun Cluster HA for SWIFTAlliance Access	SUNW.gds
Sun Cluster HA for SWIFTAlliance Gateway	SUNW.gds
Sun Cluster HA for Sybase ASE	SUNW.sybase
Sun Cluster HA for WebLogic Server	SUNW.wls
Sun Cluster HA for WebSphere MQ	SUNW.gds
Sun Cluster HA for WebSphere MQ Integrator	SUNW.gds

[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 support by Sun Cluster Security Hardening:

- Sun Cluster HA for Apache
- Sun Cluster HA for Apache Tomcat
- Sun Cluster HA for BEA WebLogic Server
- Sun Cluster HA for DHCP
- Sun Cluster HA for DNS
- Sun Cluster HA for MySQL
- Sun Cluster HA for N1 Grid Engine
- Sun Cluster HA for NFS
- Sun Cluster HA for Oracle E-Business Suite
- Sun Cluster HA for Oracle
- Sun Cluster Support for Oracle Real Application Clusters

- Sun Cluster HA for PostgreSQL
- Sun Cluster HA for Samba
- Sun Cluster HA for Siebel
- Sun Cluster HA for Solaris Containers
- Sun Cluster HA for SWIFTAlliance Access
- Sun Cluster HA for SWIFTAlliance Gateway
- Sun Cluster HA for Sun Java System Directory Server
- Sun Cluster HA for Sun Java System Message Queue
- Sun Cluster HA for Sun Java System Messaging Server
- Sun Cluster HA for Sun Java System Web Server
- Sun Cluster HA for Sybase ASE
- Sun Cluster HA for WebSphere MQ
- Sun Cluster HA for WebSphere MQ Integrator

[Top](#)

Workaround: This issue is not related to the Sun Cluster Data Service for SAP and it is addressed in SAP note 888312. Refer to SAP note 888312 to resolve this issue.

Known Issues and Bugs

The following known issues and bugs affect the operation of the Sun Cluster 3.2 2/08 release. Bugs and issues are grouped into the following categories:

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

[Top](#)

Administration

Node Panics when InfiniBand Transport Cable is Disabled (6732481)

Problem Summary: If the InfiniBand transport cable is disabled on a node that runs Sun Cluster 3.2 or 3.2 2/08, a panic can occur on the node.

Workaround: Perform one of the following actions:

- Upgrade to Sun Cluster 3.2 1/09.
- Apply the appropriate Sun Cluster 3.2 CORE patch: the latest version of Patch from My Oracle Support [126106](#) (Solaris 10) or Patch [126107](#) (Solaris 10_x86) is available on My Oracle Support.

[Top](#)

The reconf.pl Script Fails on a Sun SPARC Enterprise T5440 Server Running Sun Cluster Software in an LDOM Configuration (6749189)

Problem Summary: If you have a Sun SPARC Enterprise T5440 Server that runs Sun Cluster software in an LDOM configuration, do not run the reconf.pl script. This script can change device paths and device instances, and can damage or destroy your Sun Cluster configuration. Recovery is extremely difficult, if not impossible.

Workaround: Do not run the reconf.pl script with this configuration.

[Top](#)

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

Problem Summary: The `clnode remove -force` command should remove nodes from the metaset. The Sun Cluster System

Administration Guide for Solaris OS provides procedures for removing a node from the cluster. These procedures instruct the user to run the `metaset` command for the Solaris Volume Manager disk set removal prior to running `clnode remove`.

Workaround: If the procedures were not followed, it might be necessary to clear the stale node data from the CCR in the usual way:

1. From an active cluster node, use the `metaset` command to clear the node from the Solaris Volume Manager disk sets.
2. Run `clnode clear -force obsolete_nodename`.

[Top](#)

Entries in `Auxnodelist` Property Cause `SEGV` During Scalable Resource Creation (6494243)

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

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

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

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

[Top](#)

More Validation Checks Needed When Combining DIDs (6605101)

Problem Summary: The `/usr/cluster/bin/scdidadm` and `/usr/cluster/bin/cldevice` commands are currently unable to verify that replicated SRDF devices being combined into a single DID device are in fact replicas of each other and belong to the specified replication group.

Workaround: Take special care when combining DID devices for use with SRDF. Ensure that the DID device instances specified are replicas of each other and that they belong to the specified replication group.

[Top](#)

Sun NAS Quorum Is Offline After a Cluster Reboot (6614299)

Problem Summary: A quorum device might be offline when the cluster nodes are rebooted under the following conditions:

- An iSCSI LUN on a supported NAS appliance is configured to provide quorum to a cluster.
- That NAS appliance is present on a different subnet than the cluster nodes.

Workaround: Configure the NAS appliance on the same subnet as the cluster nodes and ensure that NAS target LUNs are statically configured targets.

[Top](#)

Node Panics Due to Loss of Quorum When iSCSI Array (STK 2510) LUN Is Used as Quorum Device (6658472)

Problem Summary: A quorum device might be offline when the cluster nodes are rebooted under the following conditions:

- An iSCSI LUN on a supported STK 2510 provides quorum to a cluster.
- That STK 2510 is present on a different subnet than the cluster nodes.

Workaround: Configure the STK 2510 on the same subnet as the cluster nodes.

[Top](#)

Interrupting `cldevice combine` Can Leave the DID Device Configuration in an Inconsistent

State (6616497)

Problem Summary: The `/usr/cluster/bin/cldevice combine` command has to update multiple pieces of information in several different places, including the DID driver's in-kernel information and multiple CCR tables. Interrupting the command while it is still in progress can leave some of these pieces of information out of date.

Workaround: If possible, avoid interrupting the `cldevice` command while it is in progress. If the command is interrupted, perform one of the following actions as applicable:

- If only the DID driver's in-kernel information was updated, and no CCR tables were modified, rebooting the cluster node on which the DID driver's information was updated returns the data to its prior state, and the `cldevice` command can be rerun.
The CCR tables that are updated are `/etc/cluster/ccr/replicated_devices` and `/etc/cluster/ccr/did_instances`. The `replicated_devices` table is updated first.
- If the `cldevice` command is interrupted before the changes to the `did_instances` table are committed, the updates to the `replicated_devices` table will have to be manually backed out.

[Top](#)

Data Services

The Validate Method Of SUNW.ScalMountPoint on a Resource Is Terminated Due to SIGSEGV Signal (6606281)

Problem Summary: Writing comments in the shared-QFS `mcf` file after the file-system information can cause the `SUNW.ScalMountPoint` methods to fail, due to improper handling by the `SUNW.ScalMountPoint` resource type. The following is an example of an `mcf` file that can lead to a problem:

```
# File system for data
Data                300    ms    Data    on    shared
#/dev/md/oracle/dsk/d30  301    md    Data    on
/dev/md/oracle/dsk/d31  302    md    Data    on
```

Workaround: Avoid writing the comments after the file-system information in the `mcf` file.

Creating an nfs resource throws an error message on s10u3 (6496879)

Problem Summary: Creating an `nfs` resource throws the following error message on `s10u3`.

```
WARNING: lockd: cannot contact statd (error 4), continuing
```

Workaround: The warning message shows up as `lockd` is being shut down while it is waiting for the `rpc` call to contact `statd`, because `statd` was not able to completely start up or is down. That `rpc` call is interrupted when `svcadm disable nlockmgr` is issued. This message is harmless and can be ignored.

Sun Cluster does not support `showmount (1M)` and `dfmounts (1M)` (6503661)

Problem Summary: Using either `dfmounts` or `showmount -a` commands to display the current NFS mounts from a node in a Sun Cluster configured as an NFS server does not result in a correct list of the active mounted resources. Instead, the list is usually empty, but can occasionally contain a few entries that may only correspond to historical client activity.

Workaround: The large number of clients in `/etc/rmtab` slows down the `mountd` startup in NFSv3, which leads to start timeouts in the cluster. Hence, `/etc/rmtab` is truncated for `mountd` to start fast. Any changes to the HA-NFS agent to patch up this problem would lead to switchover or failover issues.

Sun Cluster Data Service for SAP Web Application Server resource groups failing over when pub net is brought down on all the nodes simultaneously (6554601)

Problem Summary: Sun Cluster Data Service for SAP Web Application Server resource groups that have another resource group with strong negative affinity with them going offline. The attempt to failover the resource group with strong negative affinity

fails, but at the same time the resource group that is declared in `rg_affinities` with `--` goes to offline on the other node. When we try to bring down the resource groups the resources goes into stop failed state because they need the pub net. The resource group goes to `stop_failed` state and the resource group remains unavailable even after pub net is restored until the user intervenes.

Workaround: Comment out `/net` in `/etc/auto_master` and remove `nis` under `automount`, `hosts`, `ipnodes` in the `/etc/nsswitch.conf` file.

For SAP version 7.1, the SAP Web Application Server primary instance resource goes to **Degraded** mode after switchover of `eng-msg` resource group (6759302)

Problem Summary: In the Sun cluster environment for SAP 7.1, the SAP Web Application Server probe returns an error message with status `Degraded` after switch over of `eng-msg` resource group.

Once the switch over of `eng-msg` resource is complete, the SAP Web Application Server restarts due to a dependency on the messaging server. The restart of the SAP Web Application Server fails and returns an error message "error 503 service unavailable".

Workaround: Refer to SAP note 966416 and follow the instructions to remove all `krnlreg` entries from the profile in order to prevent deadlock situations.

For SAP version 7.1, the SAP Web Application Server go to **STOP_Failed** mode after a public network failure (6766425)

Problem Summary: In the Sun cluster environment for SAP 7.1, the SAP Web Application Server go to `STOP_Failed` mode after a public network failure. After a public network failure, the resource should failover instead of going to `STOP_Failed` mode.

Workaround: This issue is not related to the Sun Cluster Data Service for SAP and it is addressed in SAP note 888312. Refer to SAP note 888312 to resolve this issue.

[Top](#)

GUI

InPage Editing of Resource or Resource Group Properties Does Not Work in Internet Explorer (6607871)

Problem Summary: Using the InPage Editing feature to modify resource or resource group properties does not work when you are using Sun Cluster Manager in Internet Explorer. The properties fields on the resource or resource group table listing do not display as editable fields. Instead, they remain as just the display fields. No errors messages are displayed.

Workaround: Use the Edit Properties wizard at the top of the properties table listing to edit resource or resource group properties.

[Top](#)

Installation

Auto Discovery During Installation for Sun Cluster/Transport Does Not Work for `igb` Interfaces on Solaris OS x64 Platforms (6821699)

Problem Summary: During a Sun Cluster installation, auto discovery for Sun Cluster/transport does not work for `igb` interfaces on Solaris OS x64 platforms.

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

[Top](#)

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

Problem Summary: If InfiniBand is used as the cluster transport and two adapters are on each node with two ports per adapter

and a total of two switches, the `scinstall` utility's adapter autodiscovery could suggest two transport paths that use the same adapter.

Workaround: Manually specify the transport adapters on each node.

[Top](#)

Installer Exception During `pkgrm` Due to Missing `/tmp/admin.txt` (6532980)

Problem Summary: Starting the Sun Java Availability Suite installer to install Sun Cluster software and related products fails with a Java exception. The error reported on the screen is `pkgrm: ERROR: unable to open admin file /tmp/admin.txt`. The installation can fail irrespective of the type of platform either while installing Sun Cluster software alone or even when only the Shared Components are being installed or upgraded.

Workaround: Check whether the file `/tmp/admin.txt` exists. If it exists, check whether this file has both read and write permissions for `root` user. If `root` does not have these permissions, assign the permissions to `root` by using the `chmod` command. If the file `/tmp/admin.txt` does not exist, create this file with read and write permissions for `root` user with the following content:

```
mail= instance=overwrite partial=nocheck runlevel=nocheck idepend=nocheck
rdepend=nocheck space=nocheck setuid=nocheck conflict=nocheck action=nocheck
basedir=default
```

[Top](#)

Ensure That a Solaris 10 Cluster Node Is Not Forwarding Packets (6559444)

Problem Summary: The `scinstall(1M)` command creates the appropriate configuration to ensure that a new Solaris 9 or earlier cluster node does not function as a router. In Solaris 10, the method for disabling the router function changed and `scinstall` has not yet been updated. As of the Solaris 10 release, packet forwarding is disabled by default. If a Solaris 10 machine has never had forwarding enabled, this default configuration is correct and sufficient. Scalable-service response and load balancing might be impacted if packet forwarding is enabled.

Workaround: To find out if forwarding is enabled or disabled, execute the following command:

```
# /usr/sbin/routeadm
```

IPv4 routing should be enabled and IPv4 forwarding should be disabled. IPv6 routing and forwarding should both be disabled. If IPv4 forwarding was enabled before Sun Cluster software was installed, execute the following command:

```
# /usr/sbin/routeadm -d ipv4-forwarding
```

See also the `routeadm(1M)` man page.

[Top](#)

Sun Cluster Manager Might Not Show Any Data Except Cluster Name After Installation of Sun Cluster 3.2 2/08 (6621283)

Problem Summary: When the existing version of the common agent container is 2.0 (`/usr/sbin/cacaoadm -V`), it might not be properly upgraded to 2.1 during the installation of Sun Cluster 3.2 2/08 software. The following commands illustrate the problem conditions after installation of Sun Cluster 3.2 2/08 software:

```
# cacaoadm -V2.1.0# cacaoadm list-modules | grep rbaccom.sun.cacao.rbac 1.0
```

This version of the RBAC module does not work with common agent container 2.1.

Workaround: Finish the upgrade manually on each node as follows:

```
% su# /usr/sbin/cacaoadm prepare-uninstall# cacaoadm start# cacaoadm list-modules |
grep rbaccomm.sun.cacao.rbac 2.1
```

Top

SUNW.LogicalHostname methods missing from whole-root zones (6711918)

Problem Summary: When creating `SUNW.LogicalHostname` resources for whole-root zones with `ip-type=exclusive`, the command might fail because of `VALIDATE` method failure. A syslog message similar to the following is reported on some of the nodes:

```
Jun 6 22:30:45 lzkuda-2a Cluster.RGM.fed: [ID 838032 daemon.error]
...Couldn't run method tag. Error in execve: No such file or directory.
```

Workaround: Configure a file-system resource that mounts the method directory from the global zone.

```
phys-schost# zonecfg -z sczone
zonecfg:sczone> add fs
zonecfg:sczone:fs> set dir=/usr/cluster/lib/rgm
zonecfg:sczone:fs> set special=/usr/cluster/lib/rgm
zonecfg:sczone:fs> set type=lofs
zonecfg:sczone:fs> end
zonecfg:sczone> exit
```

Top

Runtime

Sun Cluster Manager Might Not Work Correctly with Sun Cluster 3.2 2/08 and Solaris 10 10/08 OS (6584336 and 6701123)

Problem Summary: If you install the Sun Cluster 3.2 2/08 software with Solaris 10 10/08 OS, the Sun Cluster installation removes the correct version of the `SUNWcacaort` package and installs the wrong version of this package. The result is that Sun Cluster Manager (formerly SunPlex Manager) might not start. See 6584336 for more information.

Workaround: After you install Sun Cluster 3.2 2/08 software, reinstall the `SUNWcacaort` package from the Solaris 10 10/08 OS distribution. You must then manually synchronize the Common Agent Container keys. See the [Sun Cluster Software Installation Guide for Solaris OS](#) for instructions on how to manually synchronize the keys.

Problem Summary: Sun Cluster Manager is unable to correctly display the information in the Tasks and Topology tabs. See 6701123 for more information.

Workaround: Perform the following commands on all cluster nodes:

```
# cd /usr/cluster/lib/SunClusterManager/WEB-INF/classes
# mv ds ds.orig
# cp -r /usr/cluster/lib/ds ds
# /usr/sbin/smcwebserver restart
```

Top

Transport Warnings on Guest LDom Cluster Nodes (6678045)

Problem Summary: On LDom guest domains that are configured as cluster nodes, messages similar to the following might display when the node is rebooted:

WARNING: Received non interrupt heartbeat on phys-schost-1:ce5 - phys-schost-4:ce5 - path timeouts are likely.

Workaround: These messages are safe to ignore.

[Top](#)

Cannot Safely Link Applications Against libstlport4 on Sun Cluster (6298053)

Problem Summary: An application that is running on a Sun Cluster configuration that is linked with `stlport4` might crash when the `nscd` daemon is not running and the application does NSS-related calls.

Workaround: On each cluster node, make modifications to the following configuration files.



Note

If you run the `sccheck` utility after removing the `cluster` switch from the `nsswitch.conf` lookup table, you might receive an error message about the missing `cluster` switch. Because you intentionally removed the `cluster` switch from the file, you can ignore this message.

Make a backup copy of the `/etc/inet/hosts` file, `/etc/netmasks` file, and `/etc/nsswitch.conf` file on each node.

1. Ensure that each private hostname in the cluster is listed in the `/etc/inet/hosts` file on each cluster node. These addresses are associated with the `clprivnet0` interface. The following example shows how to determine the IP address of a node's private hostname:

```
# ifconfig clprivnet0 | fmtclprivnet0:
flags=1009843<UP,BROADCAST,RUNNING,MULTICAST,MULTI_BCAST,PRIVATE,IPv4>mtu
1500 index 5inet 172.16.4.1 netmask fffffe00 broadcast 172.16.5.255ether
0:0:0:0:0:1
# cat > /etc/inet/hosts <<EOS172.16.0.129 clusternode1-priv-physical1
172.16.1.1 clusternode1-priv-physical2172.16.4.1 clusternode1-priv
172.16.0.130 clusternode2-priv-physical1172.16.1.2
clusternode2-priv-physical2172.16.4.2 clusternode2-privEOS
```

2. Ensure that all cluster-interconnect netmasks are listed in the `/etc/netmasks` file on all cluster nodes. You can use the following script to collect the list of netmasks:

```
# ifconfig -a | \
nawk '/flags/&&!/PRIVATE/{p=0}\
/flags/&&!/PRIVATE/{p=1}\
p==1&&$3 ~ /netmask/{print $2}' | \
while read i; do getent netmasks $i; done > /tmp/netmasks
```

Check the list of collected netmasks, for example:

```
# cat /tmp/netmasks172.16.1.1          255.255.255.128172.16.0.129
255.255.255.128172.16.4.1      255.255.254.0
```

When the list is correct, append the entries to the `netmasks` files:

```
# cat /tmp/netmasks >> /etc/netmasks
```

- 4) On each cluster node, remove the `cluster` switch from the lookups for the `hosts` and `netmasks` entries. For example:

Before modification:


```
...
hosts: cluster files files dns nisplus dns...
netmasks: cluster files nisplus...
```

After modification:

```
...
hosts: files files dns nisplus dns...
netmasks: files nisplus...
```

[Top](#)

EFI-Labeled Disks on Top of EMC PowerPath Configuration Cause Error Messages (6536409)

Problem Summary: Errors are seen when using EFI-labeled disks with EMC's PowerPath multipathing driver. These errors prevent the proper configuration of the DID device paths for these devices.

Workaround: There is no known workaround at this time. Avoid using EFI-labeled disks for devices that use EMC's PowerPath driver.

[Top](#)

Upgrade

vxiio Major Numbers Different Across the Nodes When the Root Disk Is Encapsulated (6445917)

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

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

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

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



Caution

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

[Top](#)

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

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

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

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

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

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

1. Become superuser.
2. Create a script with the following content.

```
#!/bin/ksh

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

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

3. Set the variables `PLATFORM`, `PATHNAME`, and `BASEDIR`.

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

PLATFORM

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

PATHNAME

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

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

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

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

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

BASEDIR

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

1. Run the script, once for the Sun Cluster framework packages and once for the data service packages. After the script is run, you should see the following message at the command prompt for each package:

```
Transferring pkgname package instance
```



Note

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

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

This is a harmless message and can be safely ignored.

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

[Top](#)

Sun Cluster Nodes Fail to Boot in Cluster Mode After Live Upgrade When Local Zones Configured With the Zone Root Path on File System Not `root (/)` (6599197)

Problem Summary: The Sun Cluster nodes fail to boot in cluster mode after a Live Upgrade. This failure happens only when local zones are configured on the nodes with their zone root path on a file system other than the root. The following error messages are seen during the bootstrap:

```
Configuring devices.NIS domain name is <domain name>
Loading smf(5) service descriptions: 15/26ERROR: unable to mount zones: zoneadm:
zone '<zonenumber>': "/usr/lib/fs/lofs/mount -{}{}'
o ro,nosub,nodevices /space/zones/<zonenumber>/root {}{}'
/.alt.sorce.15030/space/zones/<zonenumber>{}{}'
sorce.15030/lu/b" failed with exit code 33{}{}'
zoneadm: zone '<zonenumber>': cannot mount /.alt.sorce.15030/space/zones/<zonenumber>{}{
}{}sorce.15030/lu/b on {}{}7space/zones/<zonenumber>/root: Error 0{}{}'
zoneadm: zone '<zonenumber>': destroying snapshot: No such zone configured{}{}'
zoneadm: zone '<zonenumber>': call to zoneadmd failed{}{}'
ERROR: unable to mount zone <zonenumber> in </.alt.sorce.15030>{}{}'
ERROR: unmounting partially mounted boot environment file systems{}{}'
ERROR: umount: warning: /dev/dsk/c0t0d0s0 not in mnttab{}{}'
umount: /dev/dsk/c0t0d0s0 not mounted{}{}ERROR: cannot unmount </dev/dsk/c0t0d0s0>{
}{}ERROR: cannot mount boot environment by name <sorce.15030>{}{}'
Could not mount the last Live Upgrade Boot Environment{}{}'
Please reboot in noncluster mode(boot -x) and Repair
19/26syncing file systems... doneProgram terminate
```

Workaround: Perform the following steps:

1. Use Live Upgrade to upgrade just the operating system.
2. Reboot the nodes in noncluster mode.
3. Upgrade the shared components, Sun Cluster core packages, and Sun Cluster data services packages by using the standard upgrade procedures.
4. Boot the machine into cluster mode.

[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 2/08 Core Patch](#)
- [Removing the Sun Cluster 3.2 2/08 Core Patch](#)
- [Patch Management Tools](#)
- [My Oracle Support Online](#)
- [Sun Cluster Patch Lists](#)
- [Required Firmware](#)

If you are upgrading to Sun Cluster 3.2 2/08 software, see [Sun Cluster Upgrade Guide for Solaris OS](#). Applying a Sun Cluster 3.2 2/08 Core patch does not provide the same result as upgrading the software to the Sun Cluster 3.2 2/08 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 from My Oracle Support, [125510-02-1](#) (S9/SPARC), [125511-02-1](#) (S10/SPARC), or [125512-02-1](#) (S19/x64), you must have the -02 version of the patch installed before you can install later versions of the patch. If you do not have the -02 patch installed and want to install at least version -03, you must use the rebooting cluster method.

See the following list for examples of patching scenarios:

- If you have Sun Cluster 3.2 2/08 software using the Solaris 10 Operating System on SPARC with patch 125511-02 installed and want to install at least 125511-03, use the rebooting node or rebooting cluster method.
- If you have Sun Cluster 3.2 2/08 software using the Solaris 10 Operating System on SPARC without 125511-02 installed and want to install at least 125511-03, choose one of the following methods:
 - Use the rebooting cluster method to install 125511-03.
 - Install 125511-02 by using the rebooting node method and then install 125511-03 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 support account, contact your Oracle service representative or sales engineer, or register online at [My Oracle Support](#).

[Top](#)

Applying the Sun Cluster 3.2 2/08 Core Patch

Complete the following procedure to apply the Sun Cluster 3.2 2/08 core patch.

How to Apply the Sun Cluster 3.2 2/08 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. Register the new version of resource types `SUNW.HAStoragePlus`, `SUNW.ScalDeviceGroup`, and `SUNW.ScalMountPoint` that are being updated in this patch. Perform resource type upgrade on any existing resources of these types to the new versions.

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



Caution

If the Sun Cluster 3.2 2/08 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 2/08 core patch permanently to your cluster.

[Top](#)

Removing the Sun Cluster 3.2 2/08 Core Patch

Complete the following procedure to remove the Sun Cluster 3.2 2/08 core patch.

How to Remove the Sun Cluster 3.2 2/08 Core Patch

1. List the resource types on the cluster.

```
# clrt list
```

2. If the list returns `SUNW.HAStoragePlus:5`, `SUNW.ScalDeviceGroup:2`, or `SUNW.ScalMountPoint:2`, you must remove these resource types. For instructions on removing a resource type, see "[How to Remove a Resource Type](#)" in Sun Cluster Data Services Planning and Administration Guide for Solaris OS.
3. Reboot all nodes of the cluster into noncluster, single-user mode.
For instructions on rebooting cluster nodes into noncluster, single-user mode, see "[How to Boot a Cluster Node in Noncluster Mode](#)" in Sun Cluster System Administration Guide for Solaris OS.
4. Remove the Sun Cluster 3.2 2/08 core patch from each node on which you installed the patch.

```
# patchrm patch--id
```

5. If you are removing a Sun Cluster 3.2 core patch (patch 126106-17 or later for Solaris 10 SPARC or patch 126107-17 or

later for Solaris 10 x86) and you are returning to an earlier Sun Cluster 3.2 patch (patch 126106-16 or lower for Solaris 10 SPARC or patch 126107-16 or lower for Solaris 10 x86), you must perform an additional step on each cluster node. Before the first reboot on each cluster node, the pnm SMF manifest must be re-imported for the system to fully recover. Run the following command in the global zone:

```
# svccfg import /var/svc/manifest/system/cluster/pnm.xml
```

If you do not re-import the manifest, all data services could go into maintenance mode.

6. Reboot into cluster mode all of the nodes from which you removed the Sun Cluster 3.2 2/08 core patch. Rebooting all of the nodes from which you removed the Sun Cluster 3.2 2/08 core patch before rebooting any unaffected nodes ensures that the cluster is formed with the correct information in the CCR. If all nodes on the cluster were patched with the core patch, you can reboot the nodes into cluster mode in any order. For instructions on rebooting nodes into cluster mode, see "How to Reboot a Cluster Node" in Sun Cluster System Administration Guide for Solaris OS.
7. Reboot any remaining nodes into cluster mode.

[Top](#)

Patch Management Tools

Information about patch management options for the Solaris OS is available at the web sites for [Sun Connection Update Manager](#) and [Sun xVM Ops Center](#). Additional information for using the Solaris patch management utility, `patchadd`, is provided in Solaris Administration Guide: Basic Administration at <http://docs.sun.com/app/docs/prod/solaris>. Refer to the version of this manual that is published for the Solaris OS release that is installed on your system.

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](#)

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 2/08 third-party patch information is provided through a My Oracle Support (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 2/08 environment. To locate this Info Doc, log on to My Oracle Support. Type **Sun Cluster 3.x Third-Party Patches** in the search criteria box.

Before you install Sun Cluster 3.2 2/08 software and apply patches to a cluster component (Solaris OS, Sun Cluster software, volume manager software, data services software, or disk hardware), review each `README` file that accompanies the patches that you retrieved. All cluster nodes must have the same patch level for proper cluster operation.

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

[Top](#)

Sun Cluster Patch Lists

The [Sun Cluster Patch Klatch](#) provides a complete and up-to-date list of patches that you need to apply to the Solaris OS, to Sun Cluster software, and to other software in your cluster configuration, based on the versions of the particular software that you are using.

[Top](#)

Required Firmware

Patch for Cluster Support for StorageTek 2530 Array

The Sun StorageTek Common Array Manager (CAM) software, Version 6.0.1 provides SCSI3 PGR support for the Sun StorageTek 2530 array for up to three nodes. The CAM 6.0.1 patch is available from the Sun Download Center, or you can log into My Oracle Support and download [Patch 137294-01](#).

[Top](#)

Sun Cluster 3.2 2/08 Documentation

The Sun Cluster 3.2 2/08 user documentation set consists of the following collections:

- Sun Cluster 3.2 2/08 Software Manuals for Solaris OS
- Sun Cluster 3.2 2/08 Data Service Manuals for Solaris OS (SPARC Platform Edition)
- Sun Cluster 3.2 2/08 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 are not updated since the previous Sun Cluster 3.2 release. Their content, however, also applies to the Sun Cluster 3.2 2/08 release.

The Sun Cluster 3.2 2/08 user documentation is available in PDF and HTML format at the following web site:

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



Note

Beginning with the Sun Cluster 3.2 release, documentation for individual data services is not translated. Documentation for individual data services is available only in English.

[Top](#)

Sun Cluster 3.2 2/08 Software Manuals for Solaris OS

Part Number	Book Title
820-2562	Sun Cluster 3.2 2/08 Documentation Center
820-2554	Sun Cluster Concepts Guide for Solaris OS
820-2559	Sun Cluster Data Services Developer's Guide for Solaris OS
820-2561	Sun Cluster Data Services Planning and Administration Guide for Solaris OS
820-2560	Sun Cluster Error Messages Guide for Solaris OS
820-2553	Sun Cluster Overview for Solaris OS
820-2557	Sun Cluster Quick Reference
820-2756	Sun Cluster Reference Manual for Solaris OS
820-2555	Sun Cluster Software Installation Guide for Solaris OS
820-2556	Sun Cluster Quick Start Guide for Solaris OS
820-2270	Sun Cluster Upgrade Guide for Solaris OS
820-2558	Sun Cluster System Administration Guide for Solaris OS

[Top](#)

Sun Cluster 3.2 2/08 Data Service Manuals for Solaris OS (SPARC Platform Edition)

Part Number	Book Title
819-3056	Sun Cluster Data Service for Agfa IMPAX Guide for Solaris OS
819-2975	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
819-5415	Sun Cluster Data Service for Kerberos Guide for Solaris OS
820-3234	Sun Cluster Data Service for MaxDB Guide for Solaris OS
820-2571	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
820-2572	Sun Cluster Data Service for Oracle Application Server Guide for Solaris OS
820-2573	Sun Cluster Data Service for Oracle E-Business Suite Guide for Solaris OS
819-2980	Sun Cluster Data Service for Oracle Guide for Solaris OS
820-2574	Sun Cluster Data Service for Oracle RAC Guide for Solaris OS
819-5578	Sun Cluster Data Service for PostgreSQL Guide for Solaris OS
820-2575	Sun Cluster Data Service for Samba Guide for Solaris OS
820-2566	Sun Cluster Data Service for SAP Guide for Solaris OS
820-2567	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
820-2569	Sun Cluster Data Service for Siebel Guide for Solaris OS
820-2578	Sun Cluster Data Service for Solaris Containers Guide
820-3702	Sun Cluster Data Service for Sun Grid Engine Guide for Solaris OS
819-2988	Sun Cluster Data Service for Sun Java System Application Server Guide for Solaris OS
819-2987	Sun Cluster Data Service for Sun Java System Application Server EE (HADB) Guide for Solaris OS
819-2989	Sun Cluster Data Service for Sun Java System Message Queue Guide for Solaris OS
820-2579	Sun Cluster Data Service for Sun Java System Web Server Guide for Solaris OS
820-2576	Sun Cluster Data Service for SWIFTAlliance Access Guide for Solaris OS
820-2577	Sun Cluster Data Service for SWIFTAlliance Gateway Guide for Solaris OS
820-2570	Sun Cluster Data Service for Sybase ASE Guide for Solaris OS
820-2580	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 2/08 Data Service Manuals for Solaris OS (x86 Platform Edition)

Part Number	Book Title
819-2975	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
819-5415	Sun Cluster Data Service for Kerberos Guide for Solaris OS
820-3234	Sun Cluster Data Service for MaxDB Guide for Solaris OS
820-2571	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
820-2572	Sun Cluster Data Service for Oracle Application Server Guide for Solaris OS
819-2980	Sun Cluster Data Service for Oracle Guide for Solaris OS
820-2574	Sun Cluster Data Service for Oracle RAC Guide for Solaris OS
819-5578	Sun Cluster Data Service for PostgreSQL Guide for Solaris OS
820-2575	Sun Cluster Data Service for Samba Guide for Solaris OS
820-2566	Sun Cluster Data Service for SAP Guide for Solaris OS
820-2568	Sun Cluster Data Service for SAP Web Application Server Guide for Solaris OS
820-2578	Sun Cluster Data Service for Solaris Containers Guide
820-3702	Sun Cluster Data Service for Sun Grid Engine Guide for Solaris OS
819-2987	Sun Cluster Data Service for Sun Java System Application Server EE (HADB) Guide for Solaris OS
819-2988	Sun Cluster Data Service for Sun Java System Application Server Guide for Solaris OS
819-2989	Sun Cluster Data Service for Sun Java System Message Queue Guide for Solaris OS
820-2579	Sun Cluster Data Service for Sun Java System Web Server Guide for Solaris OS
820-2580	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
819-2995	Sun Cluster 3.1 - 3.2 With SCSI JBOD Storage Device 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
819-3020	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
819-2994	Sun Cluster 3.1 - 3.2 With Fibre Channel JBOD Storage Device Manual
817-5681	Sun Cluster 3.1 - 3.2 With SCSI JBOD Storage Device Manual for Solaris OS
819-3023	Sun Cluster 3.1 - 3.2 With Sun StorEdge T3 or T3+ Array Manual for Solaris OS
819-3019	Sun Cluster 3.1 - 3.2 With Sun StorEdge 6130 Array Manual
819-3024	Sun Cluster 3.1 - 3.2 With Network-Attached Storage Devices 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
817-0180	Sun Cluster 3.1 - 3.2 With Sun StorEdge 3310 or 3320 SCSI RAID Array Manual for Solaris OS
819-3024	Sun Cluster 3.1 - 3.2 With Network-Attached Storage Devices Manual for Solaris OS
819-3021	Sun Cluster 3.1 - 3.2 With Sun StorEdge 9900 Series Storage Device Manual for Solaris OS
819-3020	Sun Cluster 3.1 - 3.2 With Sun StorEdge 6320 System Manual for Solaris OS
819-3019	Sun Cluster 3.1 - 3.2 With Sun StorEdge 6130 Array Manual
819-3018	Sun Cluster 3.1 - 3.2 With Sun StorEdge 6120 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
819-2995	Sun Cluster 3.1 - 3.2 With SCSI JBOD Storage Device Manual for Solaris OS

Top

Documentation Issues

This section discusses errors or omissions for documentation, online help, or man pages in the Sun Cluster 3.2 2/08 release.

- [Software Installation Guide](#)
- [System Administration Guide](#)
- [Data Services Developer's Guide](#)
- [Sun Cluster Data Service for Sybase ASE Guide](#)
- [Network-Attached Storage Devices Manual](#)
- [Man Pages](#)

Top

Software Installation Guide

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

Location of Sun Cluster Plug-in for Service Provisioning System

In procedures in Chapter 3, "Establishing the Cluster", the location of documentation for the Sun Cluster Plug-in for Sun N1 Service Provisioning System is missing from the HTML version and incorrect in the PDF version. The correct location of this documentation is the following:

<http://wikis.sun.com/display/SunCluster/Sun+Cluster+Framework+Plug-In>

[Top](#)

System Administration Guide

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

Renaming a Replication Device

If the name of a replication device group changes, additional steps are required for Hitachi TrueCopy and SRDF. Follow the steps in [How to Configure DID Devices for Replication Using EMC Symmetrix Remote Data Facility \(SRDF\)](#) in Sun Cluster System Administration Guide for Solaris OS, and perform the following steps:

For TrueCopy:

If the name of the replication device group (and the corresponding global device group) changes, you must rerun the `cldevice replicate` command to update the replicated device information.

For SRDF:

If the name of the replication device group (and the corresponding global device group) changes, you must update the replicated device information by first using the `scdidadm -b` command to remove the existing information. The last step is to use the `cldevice combine` command to create a new, updated device.

Releasing SCSI-2 Reservations

When you apply patches in single-user mode with failover zones and SCSI-2 reservations exist on the disks, follow the instructions in see [How to Apply Patches in Single-User Mode with Failover Zones](#), with the following change:

After you complete Step #10, skip Step #11 and follow these instructions to determine if SCSI-2 reservations exist on the disks in the diskset.

- a. For all disks in the diskset, run `/usr/cluster/lib/sc/scsi -c disfailfast -d /dev/did/rdisk/d#s2`.
- b. If keys are listed, release them with: `/usr/cluster/lib/sc/scsi -c release -d /dev/did/rdisk/d#s2`.

Data Services Developer's Guide

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

Method Timeout Behavior Is Changed

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



Note

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

Sun Cluster Data Service for Sybase ASE Guide

This section discusses the errors and omissions in Sun Cluster Data Service for Sybase ASE Guide.

HA-Sybase support of non-global zones in SPARC architecture

In the Sun Cluster Data Service for Sybase ASE guide the information that HA-Sybase can be installed and configured on SPARC architecture is missing. Installation and configuration of HA-Sybase can be performed on non-global zones in a SPARC architecture also and this is a documentation error.

[Top](#)

Network-Attached Storage Devices Manual

The section "Requirements When Configuring Sun NAS Devices as Quorum Devices" omits the following requirement:

The Sun NAS device must be located on the same network as the cluster nodes. If a Sun NAS quorum device is not located on the same network as the cluster nodes, the quorum device is at risk of not responding at boot time within the timeout period, causing the cluster bootup to fail due to lack of quorum.

[Top](#)

Man Pages

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

[Top](#)

cldevice(1CL)

The following information for the `cldevice replicate` command is missing from the `cldevice(1M)` man page:

`replicate`

The `replicate` subcommand is not a supported method for combining DID instances with SRDF. This subcommand can be used only with TrueCopy. To combine DID instances with SRDF, use `cldevice combine`.

[Top](#)

clresource(1CL)

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

`-u`

With the `+` operand, specifies that the command operates on resources whose resource group is suspended. If you do not also specify the `-u` option when you specify the `+` operand, the command ignores all resources whose resource group is suspended. The `-u` option is valid when the `+` operand is specified with the `clear`, `disable`, `enable`, `monitor`, `set`, and `unmonitor` subcommands.

- The description of the `+` operand should state that, when used with the `clear`, `disable`, `enable`, `monitor`, `set`, or `unmonitor` subcommand, the command ignores all resources whose resource group is suspended, unless you also specify the `-u` option.
- The example provided in the definitions of the `+` and `-` operands for the `-p`, `-x`, and `-y` options are incorrect. The definitions should be as follows:

`+`

Adds a value or values to a string array value. Only the `set` subcommand accepts this operator. You can specify this operator only for the properties that accept lists of string values, for example, `Resource_dependencies`.

`-`

Deletes a value or values from a string array value. Only the `set` subcommand accepts this operator. You can specify this operator only for properties that accept lists of string values, for example `Resource_dependencies`.

[Top](#)

clresourcegroup(1CL)

- The command syntax and description for the `evacuate` subcommand incorrectly states that you can evacuate more than one node or zone in the same command invocation. Instead, you can specify only one node or zone in the `evacuate` command.
- The man page incorrectly states that the `-s` state option is available for use with the `list` subcommands. The `-s` option is currently not available in the `clresourcegroup` command for the `list` subcommand. This option is only available for the `status` subcommand. Ignore all information in the man page about the `-s` option when used with the `list`

subcommand.

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

`-u`

Specifies that the command operates on suspended resource groups, if you specify the `+` operand. If you do not also specify the `-u` option when you specify the `+` operand, the command ignores all suspended resource groups.

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

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

[Top](#)

r_properties(5)

The following new resource property is added:

Global_zone_override (boolean)

This property is allowed only for resource types that set the `Global_zone=TRUE` property in the RTR file. The setting of the `Global_zone_override` property overrides the value of the resource type property `Global_zone` for the particular resource.

Setting the `Global_zone_override` property to `FALSE` forces the resource methods to execute in the non-global zone in which the resource group is configured, rather than always executing in the global zone as they usually would when the `Global_zone` property is set to `TRUE`.

This property is optional if a default value is specified in the RTR file.

If the `Tunable` attribute is not specified in the RTR file, the `Tunable` value for the property is `At_creation`. You can set the `Tunable` attribute in the RTR file to `At_creation`, `When_disabled`, or `Anytime`.

Use caution when you set the `Tunable` attribute to `Anytime` in the RTR file. Changes to the `Global_zone_override` property take effect immediately, even if the resource is online. For example, suppose that the `Global_zone_override` tunability is set to `ANYTIME` and the `Global_zone_override` property is currently set to `FALSE` on a resource that is configured in a non-global zone. When the resource is switched online, the starting methods are executed in the non-global zone. If the `Global_zone_override` property is then set to `TRUE` and the resource is switched offline, the stopping methods are executed in the global zone. Method code must be able to deal with this possibility. If it cannot, then you must set the `Tunable` attribute to `When_disabled` or `At_creation` instead.

Category	Conditional/Optional
Default	TRUE
Tunable	At creation

[Top](#)

rt_properties(5)

API_version

In the description of the `API_version` property, the date shown for the release (12/07) is incorrect. The date should be 2/08, as shown next.

```
3.2 2/08
```

```
| 8
```

[Top](#)

Global_zone

The following information describes the effect that the new `Global_zone_override` resource property has on the `Global_zone` resource-type property:

A resource type that declares `Global_zone=TRUE` might also declare the `Global_zone_override` resource property. In that case, the value of the `Global_zone_override` property supersedes the value of the `Global_zone` property for that resource.

[Top](#)

`scdidadm(1M)`

The following information about the new `-b` option is missing from the `scdidadm(1M)` man page:

`-b`

Returns a replicated DID instance to its prior state of being two separate DID instances.

Example 5.5 Undo Replicated DID Device Configuration

```
# scdidadm -b 10
```

[Top](#)

Resolved Issues

This section discusses Change Requests or documentation errors or omissions that were resolved after the Sun Cluster 3.2 2/08 release.

InfiniBand Host Channel Adapter (HCA) Support (6599044)

The Sun SPARC Enterprise T5120 and T5220 servers that use Sun Cluster 3.2 2/08 software no longer require CORE Patch 126106 (Solaris 10) or Patch 126107 (Solaris 10_x86) for InfiniBand configurations.

[Top](#)

Private Network Interfaces Using the `nxge` Driver Can Cause a Solaris OS Panic (6525548 and 6726165)

Problem Summary: Private network interfaces that use the `nxge(7D)` driver can cause a panic in the Solaris OS (6525548 and 6726165). A cluster node might panic if the network adapter meets the following criteria:

1. Your private network adapter uses the `nxge` driver.
2. A change in network adapter state follows a change in cluster configuration.

A stack trace of the panic might look like this stack trace (6525548):

```
_1cGcursorHgetdata6Mlpv_l_+0xb0(3, d0, 7b28d0dc, 60008f6d9d8, c, 4)
_1cNMarshalStreamJget_bytes6Mpvlb_v_+0x30(60008f6d9b0, 2a102d05784, 0, 0, 16, 60008f6d9d8)
_1cTckpt_handler_serverHprocess6MrnHservicev_+0x28(60005aed740, 2a102d05858, 0, 703d88a0, 70578000, 0)
_1cJckpt_elemHexecute6M_v_+0x50(60006c4b3c0, 0, 703aa800, 60008f6d980, 60006c4b3e0, 2a102d05858)
_1cTthreadpool_worker_tVdeferred_task_handler6M_v_+0x120(703d9078, 60004b071c0, 60006c4b3c0, 703d8820, 1, 703d8820)
_1cKthreadpoolOthread_handler6FpnTthreadpool_worker+t{+}v+0x24(60004b071c0, 1883800, 703fac00, 703fac00, 0, 0)
cllwpwrapper+0xf8(7b357cac, 0, 70422c00, 70422c00, 0, 0)
thread_start+4(2a102d05b70, 18, 0, 0, 0, 0)
```

A demangled stack trace of the panic might look like this stack trace (6525548):

```

unsigned cursor::getdata+0xb0(3, d0, 7b28d0dc, 60008f6d9d8, c, 4)
void MarshalStream::get_bytes+0x30(60008f6d9b0, 2a102d05784, 0, 0, 16, 60008f6d9d8)
void ckpt_handler_server::process+0x28(60005aed740, 2a102d05858, 0, 703d88a0, 70578000, 0)
void ckpt_elem::execute+0x50(60006c4b3c0, 0, 703aa800, 60008f6d980, 60006c4b3e0, 2a102d05858)
void threadpool_worker_t::deferred_task_handler+0x120(703d9078, 60004b071c0, 60006c4b3c0, 703d8820, 1,
703d8820)
void threadpool::thread_handler+0x24(60004b071c0, 1883800, 703fac00, 703fac00, 0, 0)
cllpwrapper+0xf8(7b357cac, 0, 70422c00, 70422c00, 0, 0)
thread_start+4(2a102d05b70, 18, 0, 0, 0, 0)

```

A stack trace of the panic might look like this stack trace (6726165):

```

> ::stack
_1cHserviceKget_object6MpnSMarshalInfo_objectpncFORBAGObject_+0x3c(2a10461b860, 706c1bc8, 4, 2e0,
60017b41e08, 5c)
_1cTckpt_handler_serverHprocess6MrnHservicev+0x58(60019c0d040, 2a10461b860, 5c5c5c5c, 7071f3d8, 0, 0)
_1cjkpt_elemHexecute6M_v+0x40(6001a9f4dc0, 60010b2ac40, 2a10461b860, 706cbc38, 706cb000, 706cb)
_1cTthreadpool_worker_tvdeferred_task_handler6M_v+0x114(7071fc10, 60017dec580, 6001a9f4dc0, 7071f328, 1,
7071f328)
_1cKthreadpoolOthread_handler6FpnTthreadpool_worker_tv+0x1c(60017dec580, 18cec00, 60017e03c40, 0,
706ee000, 706ee)
cllpwrapper+0xc4(2a10461bb70, 7b342de4, 0, 0, 70702000, 70702)
thread_start+4(2a10461bb70, 18, 0, 0, 0, 0)

```

A demangled stack trace might look like this stack trace (6726165):

```

> ::stack
CORBA::Object*service::get_object+0x3c(2a10461b860, 706c1bc8, 4, 2e0, 60017b41e08, 5c)
void ckpt_handler_server::process+0x58(60019c0d040, 2a10461b860, 5c5c5c5c, 7071f3d8, 0, 0)
void ckpt_elem::execute+0x40(6001a9f4dc0, 60010b2ac40, 2a10461b860, 706cbc38, 706cb000, 706cb)
void threadpool_worker_t::deferred_task_handler+0x114(7071fc10, 60017dec580, 6001a9f4dc0, 7071f328, 1,
7071f328)
void threadpool::thread_handler+0x1c(60017dec580, 18cec00, 60017e03c40, 0, 706ee000, 706ee)
cllpwrapper+0xc4(2a10461bb70, 7b342de4, 0, 0, 70702000, 70702)
thread_start+4(2a10461bb70, 18, 0, 0, 0, 0)
>

```

To determine whether you are using the `nxge` driver, type one of the following commands:

- If you are using Sun Cluster 3.2, on a cluster node, type `clinterconnect status`.
- If you are using Sun Cluster 3.1, on a cluster node, type `scstat -w`, as shown in the following example.

```

# scstat -w -- Cluster Transport Paths -- Endpoint
Endpoint          Status          Endpoint

Transport path: phys-host-1:nxge1 phys-host-2:nxge1 Path online
Transport path: phys-host-1:nxge0 phys-host-2:nxge0 Path online#

```



Note

You can also use the `modinfo` or the `ifconfig` command to determine if the `nxge` driver is used on a node. For example, you can type `modinfo -c | grep nxge`

Examples of a change in adapter state that follows a change in cluster configuration are:

- The cluster switches over and the private adapter fails.
- The private adapter is disabled.

Solution: For SPARC, apply Patch 138048-03 or later to resolve both the Solaris OS panic and the Sun Cluster software issue. For x86, apply Patch 138049-03 or later.

Workaround: You can perform the following workaround until you apply the appropriate patch:

1. On each node that uses an `nxge` private adapter, add the following entry to the `/etc/system` file.

```
set nxge:nxge_rx_threshold_hi=0
```

2. Reboot.

[Top](#)

Required Minimum Solaris Kernel Patch for UltraSPARC T2 and UltraSPARC T2 Plus Servers

Problem Summary: The Sun Cluster software requires Solaris kernel Patch 137111-01 or later to avoid a node panic on all UltraSPARC T2 and UltraSPARC T2 Plus-based servers. Without this patch, a panic can occur when a node in the cluster attempts to mount a Global File System (pxfs).

A stack trace of the panic is similar to the following:

```
> ::stack
fop_ioctl+0x20(6001bafaac0, 20006657, 2a1044d900c, 200000, 60010803e48, 1313624)
_1cRpxfs_namespace_v1Rfs_dependent_implMkernel_ioctl6MpnDvfs_il_i+0x110(18f0, 706caed8, 20006657,
2a1044d900c, 70400, 1800)
_1cRpxfs_namespace_v1Sufs_dependent_implSconvert_to_primary6Mpn0AFfs_iii+0x100(6001597ff30,
6001198d510, 0, 30004878548, 3000555c3c4, 0)
_1cRpxfs_namespace_v1Ffs_iiSconvert_to_primary6Mb_i+0x214(6001198d510, 0, 1, 7a6aac2c, 7088c860,
6001198d570)
_1cRpxfs_namespace_v1Qrepl_pxfs_serverObecome_primary6MrknL_string_seq_rnFCORBAEnvironmentv
+0x950(6001bad5c00, 6001bad5c00, 2a1044d9828, 6001198d500, 6001913c648, 6001b1e8508)
_1cRgeneric_repl_provObecome_primary6MrknL_string_seq_rnFCORBAEnvironmentv+0xe4(60019e66540,
2a1044d9500, 2a1044d9828, 70400, 1, 2a1044d9358)
_1cNrma_prov_implObecome_primary6MrknL_string_seq_rnFCORBAEnvironmentv+0x78(60017b69910, 7,
2a1044d9828, 7071dc08, 7aa2c240, 7080e418)
_1cbX_replica_int_rma_repl_prov_become_primary_receive6FpvrnHservicev+0x68(60017b69918, 2a1044d97c8,
706cac18, 2a1044d9828, 706ca000, 2a1044d9500)
_1cOremote_handlerUhandle_incoming_call6MrnHservicev+0x10c(60017b69920, 2a1044d97c8, 7b2f1920,
60017b69918, 70731348, 7b318e70) 600173c5a00, 2a1044d9828)
_1cGrxdoorNhandle_twoway6Fpnjrecstreamv+0xf8(60014ca5380, 70400, 60014ca53b0, b9, 60014ca5450,
2a1044d97c8)
_1cTthreadpool_worker_tvdeferred_task_handler6M_v+0x114(7080e418, 60017de10a8, 60014ca5380,
60017936f98, 2, 60017936f98)
_1cKthreadpoolOthread_handler6FpnTthreadpool_worker_tv+0x1c(60017de10a8, 18cec00, 60017dd96a0, 0,
706ee000, 706ee)
cllwpwrapper+0xc4(2a1044d9b70, 7b330de4, 0, 0, 70702000, 70702)
thread_start+4(2a1044d9b70, 18, 0, 0, 0, 0)
```

A demangled stack trace of the panic is similar to the following:

```
> $G
C++ symbol demangling enabled
> ::stack
fop_ioctl+0x20(6001bafaac0, 20006657, 2a1044d900c, 200000, 60010803e48, 1313624)
int pxfs_namespace_v1::fs_dependent_impl::kernel_ioctl+0x110(18f0, 706caed8, 20006657, 2a1044d900c, 70400,
1800)
int pxfs_namespace_v1::ufs_dependent_impl::convert_to_primary+0x100(6001597ff30, 6001198d510, 0,
30004878548, 3000555c3c4, 0)
int pxfs_namespace_v1::fs_ii::convert_to_primary+0x214(6001198d510, 0, 1, 7a6aac2c, 7088c860, 6001198d570)
void pxfs_namespace_v1::repl_pxfs_server::become_primary+0x950(6001bad5c00, 6001bad5c00, 2a1044d9828,
6001198d500, 6001913c648, 6001b1e8508)
void generic_repl_prov::become_primary+0xe4(60019e66540, 2a1044d9500, 2a1044d9828, 70400, 1, 2a1044d9358)
void rma_prov_impl::become_primary+0x78(60017b69910, 7, 2a1044d9828, 7071dc08, 7aa2c240, 7080e418)
void _replica_int_rma_repl_prov_become_primary_receive+0x68(60017b69918, 2a1044d97c8, 706cac18,
2a1044d9828, 706ca000, 2a1044d9500)
void remote_handler::handle_incoming_call+0x10c(60017b69920, 2a1044d97c8, 7b2f1920, 60017b69918,
70731348, 7b318e70)
void rxdoor::handle_request_common+0x3f4(60014ca5450, 2a1044d97c8, 2a1044d981c, 8, 600173c5a00,
2a1044d9828)
void rxdoor::handle_twoway+0xf8(60014ca5380, 70400, 60014ca53b0, b9, 60014ca5450, 2a1044d97c8)
void threadpool_worker_t::deferred_task_handler+0x114(7080e418, 60017de10a8, 60014ca5380, 60017936f98, 2,
60017936f98)
void threadpool::thread_handler+0x1c(60017de10a8, 18cec00, 60017dd96a0, 0, 706ee000, 706ee)
cllwpwrapper+0xc4(2a1044d9b70, 7b330de4, 0, 0, 70702000, 70702)
thread_start+4(2a1044d9b70, 18, 0, 0, 0, 0)
```

Solution: Apply Patch 137111-01 to eliminate the cause of the Solaris kernel panic.

[Top](#)