

Sun Cluster 3.2 Release Notes Supplement

3.2 Release Notes Supplement

820-3796

Release Notes Supplement

Sun Cluster 3.2 Release Notes Supplement

820-3796

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

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

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

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

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

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

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

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

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

Sun Cluster 3.2 Release Notes Supplement

Sun Cluster 3.2 Release Notes Supplement

This page supplements the standard user documentation, including the [Sun Cluster 3.2 Release Notes for Solaris OS](#) that shipped with the Sun™ Cluster 3.2 product. These "online release notes" provide the most current information on the Sun Cluster 3.2 product. This page includes the following information.

- [Revision Record](#)
- [New Features](#)
- [Restrictions and Requirements](#)
- [Known Problems](#)
- [Known Documentation Problems](#)
- [Resolved Issues](#)

[Top](#)

Revision Record

The following tables list the information contained in this page and provides the revision date for this information.

Sun Cluster 3.2 Release Notes Supplement Revision Record 2008

Revision Date	New Information
After January 2008	Go to History (under the Tools icon) for the list of changes made to this document
January 2008	Support Limitation for Exclusive-IP Zones

Sun Cluster 3.2 Release Notes Supplement Revision Record 2007

Revision Date	New Information
October 2007	Required Steps if Solaris Secure by Default Option Is Enabled Need Documentation (6558275)
	Support for Live Upgrade of VxVM 5.0 With Required Patch
	Support for Sun Cluster 3.1 Agent Software
	You Can Now Configure e1000g Cards as Cluster Interconnect Using Back-to-Back Cabling
	Private Network Interfaces Using the <code>nxge</code> Driver Might Panic Solaris (6525548)
September 2007	Support for Sun StorageTek 9900 Systems as Quorum Devices in Campus Clusters
August 2007	New requirements for the Sun Fire™ V215 and Sun Fire V245 Platforms That Are Running Solaris 9

[Top](#)

New Features

In addition to features documented in the [Sun Cluster 3.2 Release Notes for Solaris OS](#), this release now includes support for the following features.

[Top](#)

Support for Sun Cluster 3.2 and Solaris 10 10/08 OS

Sun Cluster 3.2 software with the latest patches is now supported on the Solaris 10 10/08 OS.

[Top](#)

Support for Sun Cluster 3.1 Agent Software

All the Sun Cluster 3.1 agents are supported in the Sun Cluster 3.2 release. If you upgrade Sun Cluster 3.1 software to Sun Cluster 3.2 software, we recommend that you upgrade all agents to Sun Cluster 3.2 to utilize any new features and bug fixes in the agent software. If you upgrade the application software, you must apply the latest agent patches to make the new version of the application highly available on Sun Cluster.

[Top](#)

Support of Sun StorageTek 9900 Systems as Quorum Devices in Campus Clusters

Sun StorageTek 9900 systems are now supported as quorum devices in campus cluster configurations. The following patches are required:

- [126105-01 Sun Cluster 3.2: CORE patch for Solaris 9](#)
- [126106-01 Sun Cluster 3.2: CORE patch for Solaris 10](#)
- [126107-01 Sun Cluster 3.2: CORE patch for Solaris 10_x86](#)

[Top](#)

Restrictions and Requirements

This section describes restrictions and requirements for Sun Cluster 3.2.

[Top](#)

Support Limitations for Exclusive-IP Zones

On versions of the Solaris OS that include the exclusive-IP feature for non-global zones, all non-global zones that are managed by Sun Cluster software must be configured as shared-IP zones. The exclusive-IP zones feature, which was introduced in the Solaris 10 11/06 release, is not currently supported by Sun Cluster software. However, you can create exclusive-IP zones on a cluster node as long as those zones remain outside Sun Cluster management. No additional steps are required to configure a zone as a shared-IP zone, as this is the default zone configuration on versions of the Solaris OS that feature exclusive-IP zones.

[Top](#)

Requirements to Use Live Upgrade with Veritas Volume Manager 5.0

You can use Live Upgrade to upgrade to Veritas Volume Manager (VxVM) 5.0 as long as you also install a required VxVM patch:

1. Upgrade to VxVM 5.0 in the alternate boot environment (BE).
2. Install the patch 122058-06 in the alternate BE. This patch is released along with the VxVM 5.0 MP1 patch set.

```
# patchadd -R /BE-mount-point 122058-06
```

[Top](#)

Requirements for Operating Shared Fibre Channel Storage Attached to Sun Fire V215 and Sun Fire V245 Platforms That Are Running Solaris 9 9/05 HW

Sun Fire V215 and Sun Fire V245 platforms that are running Sun Cluster 3.1 or Sun Cluster 3.2 on Solaris 9 9/05 HW require that you apply at least version 10 of kernel patch [122300-13 SunOS 5.9: Kernel Patch](#) and at least version 4.4.13 of [Sun StorEdge™ SAN Foundation Kit](#) for proper operation with shared Fibre Channel storage. Both URLs (for the kernel patch and SAN Foundation Kit) require you to provide a user ID and password.

[Top](#)

You Can Now Configure e1000g Cards as Cluster Interconnect Using Back-to-Back Cabling

Previously, you could not use switches for private interconnects in the following hardware configurations, even if the Sun Cluster configuration included two nodes:

- Sun Fire T2000 onboard Ethernet ports
- Sun Netra T2000 onboard Ethernet ports
- Sun PCI-Express Dual Gigabit Ethernet UTP Adapter, (X)7280A-2

Now, if you have installed [Solaris Operating System 10 Kernel Update patch 118833-30](#) or later versions of this patch, you can use switches for private interconnects in the preceding two-node hardware configurations.

[Top](#)

Known Problems

In addition to known problems that are documented in [Sun Cluster 3.2 Release Notes for Solaris OS](#), the following known problems affect the operation of the Sun Cluster 3.2 release.

[Top](#)

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 [126106](#) (Solaris 10) or Patch [126107](#) (Solaris 10_x86) is available on SunSolve.

[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](#)

InfiniBand and Sun Cluster Requirements for T5120 and T5220 (6599044)

Problem Summary: Sun SPARC Enterprise T5120 and T5220 servers require Sun Cluster 3.2 2/08 software or later when configured with InfiniBand for Sun Cluster 3.2. If you are running Sun Cluster 3.1 software, the minimum version is the standard InfiniBand requirement of Sun Cluster 3.1 8/05. An InfiniBand configuration requires Solaris 10 1/06 or later.

If you are running Sun Cluster 3.2 GA software, apply the appropriate CORE patch: the latest version of Patch [126106](#) (Solaris 10) or Patch [126107](#) (Solaris 10_x86) is available on SunSolve.

[Top](#)

Required Steps if Solaris Secure by Default Option is Enabled Need Documentation (6558275)

Problem Summary: During the installation of the Solaris 10 11/06 or Solaris 10 8/07 OS, if you choose not to enable network services for remote clients, a restricted network profile is used that disables external access for certain network services. The restricted services include the following services that affect cluster functionality:

- The RPC communication service, which is required for cluster communication
- The Sun Java Web Console service, which is required to use the Sun Cluster Manager GUI

Workaround: The following steps restore Solaris functionality that is used by the Sun Cluster framework but which is prevented if a restricted network profile is used.

1. Perform the following commands to restore external access to RPC communication.

```
# svccfg
svc:> select network/rpc/bind
svc:/network/rpc/bind> setprop config/local_only=false
svc:/network/rpc/bind> quit
# svcadm refresh network/rpc/bind:default
# svcprop network/rpc/bind:default | grep local_only
```

The output of the last command should show that the local_only property is now set to false.

2. (Optional) Perform the following commands to restore external access to Sun Java Web Console.

```
# svccfg
svc:> select system/webconsole
svc:/system/webconsole> setprop options/tcp_listen=true
svc:/system/webconsole> quit
# /usr/sbin/smcwebserver restart
# netstat -a | grep 6789
```

The output of the last command should return an entry for 6789, which is the port number that is used to connect to Sun Java Web Console.

[Top](#)

CheckS6708638.java contains typos in user visible text (6805325)

Problem Summary: During scinstall, when responding "yes" to the question:

Do you want to run cluster check (yes/no) [yes]?

Users will see a possibly confusing message as follows:

```
l82bb019: S6708638.... starting: Node does has insufficient physical memory.
l82bb019: S6708638 passed
```

Workaround: This message is harmless and can be safely ignored as long as the check is indicated as "passed"

[Top](#)

Known Documentation Problems

This section discusses documentation errors you might encounter and steps to correct these problems. This information is in addition to known documentation problems that are documented in the [Sun Cluster 3.2 Release Notes for Solaris OS](#).

Incorrect Listing of HA-NetBackup Manual in the Data Services Collection

In the listing of [data-service manuals](#) for SPARC based platforms, the Sun Cluster Data Service for NetBackup Guide is erroneously listed. This data service is supported on Sun Cluster 3.1 software but not on Sun Cluster 3.2 software. Contact Symantec to determine the availability of a data service for Veritas NetBackup that is supported on Sun Cluster 3.2 software.

Locations of Documentation for Versions of the Data Service for Sun Java System Directory Server

Procedures for the version of Sun Cluster HA for Sun Java™ System Directory Server that uses Sun Java System Directory Server 5.0 and 5.1 are located in the Sun Cluster 3.1 Data Service for Sun ONE Directory Server. Beginning with Version 5.2 of Sun ONE Directory Server, see the Sun ONE Directory Server or Sun Java System Directory Server installation documentation.

Obsolete `SUNW.rac_hwraid` Resource Type Listed in Release Notes Table of Resource Types

In the "Supported Products" section of the Sun Cluster 3.2 Release Notes, the table of data services and their resource types lists the `SUNW.rac_hwraid` as one of the resource types available for Oracle RAC. This resource type was made obsolete as of the 3.2 release and should not be included in this list of resource types.

Resolved Issues

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_I_+0xb0(3, d0, 7b28d0dc, 60008f6d9d8, c, 4)
_1cNMarshalStreamJget_bytes6MpvIb_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)
cllwpwrapper+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_object6MpnSMarshallInfo_objectpnFCORBAGObject_+0x3c(2a10461b860, 706c1bc8, 4, 2e0,
60017b41e08, 5c)
_1cTckpt_handler_serverHprocess6MrnHservicev+0x58(60019c0d040, 2a10461b860, 5c5c5c5c, 7071f3d8, 0, 0)
_1cjckpt_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)
cllwpwrapper+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_tv::deferred_task_handler+0x114(7071fc10, 60017dec580, 6001a9f4dc0, 7071f328, 1,
7071f328)
void threadpool::thread_handler+0x1c(60017dec580, 18cec00, 60017e03c40, 0, 706ee000, 706ee)
cllwpwrapper+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   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)
__1cGrxdoorVhandle_request_common6FrnHID_node_rnHservice_pnSrxdoor_invo_header_C_v
+0x3f4(60014ca5450, 2a1044d97c8, 2a1044d981c, 8, 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)
cllpwrapper+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](#)