

Oracle® Solaris Cluster 3.3 3/13 Release Notes

Copyright © 2000, 2013, 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 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 END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

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 that 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 of The Open Group.

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.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée d'The Open Group.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.

Contents

Preface	7
Oracle Solaris Cluster 3.3 3/13 Release Notes	11
What's New in the Oracle Solaris Cluster 3.3 3/13 Software	11
Support for Oracle 12c	12
Qualification of VxVM and VxFS	12
New clsetup Wizards to Create a Zone Cluster	12
Resource Dependencies Can Be Defined on a Per-Node Basis	13
Support for Kernel Cage Dynamic Reconfiguration (DR)	13
Cluster Security Framework Is Enhanced	13
Faster Failure Detection and Response by Storage Monitors	13
New clsetup Wizard to Configure the Oracle PeopleSoft Application Server Data Service	14
New clsetup Wizard to Configure the Oracle WebLogic Server Data Service	14
New Data Service for Oracle External Proxy	14
New Data Service for SAP NetWeaver	14
New Data Service for Oracle PeopleSoft Enterprise Process Scheduler	15
New Data Service for Oracle Web Tier	15
Support for Sun ZFS Storage Appliance Data Replication With Geographic Edition	15
Restrictions	15
Oracle ACFS as a Cluster File System	15
Topology Pages in Oracle Solaris Cluster Manager	16
Commands Modified in This Release	16
Compatibility Issues	16
Oracle ASM With Solaris Volume Manager Mirrored Logical Volumes	16
ACFS Proxy Resource Goes Offline During Installation of Oracle Binaries (16581850)	17
Tavor InfiniBand (IB) Transport Cluster Nodes Panic: mutex_enter: bad mutex, lp (16372022)	17

Solaris Upgrade to Oracle Solaris 10 8/11 Overwrites the hostname . <i>adp</i> Contents (15882549)	17
Solaris Volume Manager GUI	18
Accessibility Information	18
Supported Products	18
Data Services	18
File Systems	19
Memory Requirements	19
Oracle Solaris Operating System	19
Oracle VM Server for SPARC	20
Availability Suite	20
Volume Managers	20
Features Not Supported in This Release	20
Sun Management Center	20
Network Appliance NAS Devices	20
Socket Direct Protocol (SDP)	20
Product Localization	21
Known Issues and Bugs	21
Administration	22
Data Services	26
Geographic Edition	31
Installation	32
Localization	34
Runtime	34
Patches and Required Firmware Levels	35
Applying an Oracle Solaris Cluster 3.3 3/13 Core Patch	36
Removing an Oracle Solaris Cluster 3.3 3/13 Core Patch	36
Patch Management Tools	37
Patch for Cluster Support for Sun StorageTek 2530 Array	37
My Oracle Support	38
Oracle Solaris Cluster 3.3 3/13 Documentation Set	38
Documentation Addendum	38
Software Installation Guide	38
HA for Oracle Guide	39
HA for Oracle RAC Guide	39
HA for Oracle Web Tier Guide	42

Man Pages 43

Preface

This document lists product features, requirements, and restrictions for Oracle Solaris Cluster 3.3 3/13 software. It also describes open bugs and other known problems.

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

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

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

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

Using UNIX Commands

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

See one or more of the following for this information:

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

Typographic Conventions

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

TABLE P-1 Typographic Conventions

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

Shell Prompts in Command Examples

The following table shows UNIX system prompts and superuser prompts for shells that are included in the Oracle Solaris OS. In command examples, the shell prompt indicates whether the command should be executed by a regular user or a user with privileges.

TABLE P-2 Shell Prompts

Shell	Prompt
Bash shell, Korn shell, and Bourne shell	\$
Bash shell, Korn shell, and Bourne shell for superuser	#
C shell	machine_name%
C shell for superuser	machine_name#

Related Documentation

Information about related Oracle Solaris Cluster topics is available in the documentation that is listed in the following table. All Oracle Solaris Cluster documentation is available at <http://www.oracle.com/technetwork/indexes/documentation/>.

Topic	Documentation
Concepts	<i>Oracle Solaris Cluster Concepts Guide</i>
Hardware installation and administration	<i>Oracle Solaris Cluster 3.3 3/13 Hardware Administration Manual</i> and individual hardware administration guides
Software installation	<i>Oracle Solaris Cluster Software Installation Guide</i>
Data service installation and administration	<i>Oracle Solaris Cluster Data Services Planning and Administration Guide</i> and individual data service guides
Data service development	<i>Oracle Solaris Cluster Data Services Developer's Guide</i>
System administration	<i>Oracle Solaris Cluster System Administration Guide</i> <i>Oracle Solaris Cluster Quick Reference</i>
Software upgrade	<i>Oracle Solaris Cluster Upgrade Guide</i>
Error messages	<i>Oracle Solaris Cluster Error Messages Guide</i>
Command and function references	<i>Oracle Solaris Cluster Reference Manual</i> <i>Oracle Solaris Cluster Data Services Reference Manual</i>

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Getting Help

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

- Your name and email address
- Your company name, address, and phone number
- The model and serial numbers of your systems
- The release number of the operating environment, for example Solaris 10

- The release number of Oracle Solaris Cluster software, for example, Oracle Solaris Cluster 3.3 3/13

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

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

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

Oracle Solaris Cluster 3.3 3/13 Release Notes

This document provides the following information for Oracle Solaris Cluster 3.3 3/13 software.

- “What's New in the Oracle Solaris Cluster 3.3 3/13 Software” on page 11
- “Restrictions” on page 15
- “Commands Modified in This Release” on page 16
- “Compatibility Issues” on page 16
- “Accessibility Information” on page 18
- “Supported Products” on page 18
- “Features Not Supported in This Release” on page 20
- “Product Localization” on page 21
- “Known Issues and Bugs” on page 21
- “Patches and Required Firmware Levels” on page 35
- “Oracle Solaris Cluster 3.3 3/13 Documentation Set” on page 38
- “Documentation Addendum” on page 38

What's New in the Oracle Solaris Cluster 3.3 3/13 Software

This section provides information related to new features, functionality, and supported products in the Oracle Solaris Cluster 3.3 3/13 software.

- “Support for Oracle 12c” on page 12
- “Qualification of VxVM and VxFS” on page 12
- “New c1setup Wizards to Create a Zone Cluster” on page 12
- “Resource Dependencies Can Be Defined on a Per-Node Basis” on page 13
- “Support for Kernel Cage Dynamic Reconfiguration (DR)” on page 13
- “Cluster Security Framework Is Enhanced” on page 13
- “Faster Failure Detection and Response by Storage Monitors” on page 13
- “New c1setup Wizard to Configure the Oracle PeopleSoft Application Server Data Service” on page 14
- “New c1setup Wizard to Configure the Oracle WebLogic Server Data Service” on page 14
- “New Data Service for Oracle External Proxy” on page 14

- “New Data Service for SAP NetWeaver” on page 14
- “New Data Service for Oracle PeopleSoft Enterprise Process Scheduler” on page 15
- “New Data Service for Oracle Web Tier” on page 15
- “Support for Sun ZFS Storage Appliance Data Replication With Geographic Edition” on page 15

Support for Oracle 12c

You can now install and configure the Oracle 12c software to work with the HA for Oracle data service. Internet Protocol version 6 (IPv6) is not supported with Oracle 12c software. See the *Oracle Solaris Cluster Data Service for Oracle Guide* for more information.

You can also install and configure Oracle 12c to work with Oracle Real Application Cluster (RAC) software, in addition to 10g release 2 and 11g. These resource types enable Oracle Clusterware and Oracle Solaris Cluster software to interoperate. Internet Protocol version 6 (IPv6) is not supported on Oracle 12c with Oracle RAC software. See “How to Enable Oracle Solaris Cluster and Oracle Clusterware 10g Release 2, 11g, or 12c to Interoperate” in *Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide* or “Creating Resources for Interoperation With Oracle 10g, 11g, or 12c by Using Oracle Solaris Cluster Maintenance Commands” in *Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide*.

Qualification of VxVM and VxFS

Veritas Volume Manager (VxVM) and Veritas File System (VxFS) delivered as part of Veritas Storage Foundation 5.1 are now qualified for use in an Oracle Solaris Cluster 3.3 3/13 configuration. Version 5.1 requires SP1 RP3. For updated information, check the My Oracle Support Oracle Solaris Cluster Information Center at <https://support.oracle.com/rs?type=doc&id=1560789.2>. Follow procedures published in the Oracle Solaris Cluster 3.3 5/11 documentation library at http://download.oracle.com/docs/cd/E18728_01/index.html.

New clsetup Wizards to Create a Zone Cluster

The `clsetup` utility now enables you to create a zone cluster and configure network addresses, storage devices, a file system, or a ZFS storage pool. For details, see “Configuring a Zone Cluster” in *Oracle Solaris Cluster Software Installation Guide*.

Resource Dependencies Can Be Defined on a Per-Node Basis

You can now specify per-node resource dependencies, which might differ for each per-node instance of a resource. Per-node instances are instances of the resource that are online simultaneously (in a multi-mastered resource group) or disjointedly in time (in a failover resource group) on different nodes. For details, see [“How to Change Resource Dependency Properties”](#) in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

Support for Kernel Cage Dynamic Reconfiguration (DR)

Kernel cage memory DR (the removal of system boards containing system memory) is now supported. Support is currently limited to Oracle's SPARC Enterprise M8000 and M9000 series servers. For details, see [“Kernel Cage DR Recovery”](#) in *Oracle Solaris Cluster 3.3 3/13 Hardware Administration Manual*.

Cluster Security Framework Is Enhanced

New facilities were added to assist in running application programs as nonroot users, securely storing and retrieving private strings such as passwords, and performing additional runtime checking of ownership and permissions for Resource Group Manager (RGM) resource type callback methods. For details, see [Chapter 13, “Security for Data Services,”](#) in *Oracle Solaris Cluster Data Services Developer's Guide*.

Faster Failure Detection and Response by Storage Monitors

The following enhancements provide quicker detection and response to storage failures:

- A new extension property, `RebootOnFailure`, is added to the `SUNW.HASStoragePlus`, `SUNW.ScalDeviceGroup`, and `SUNW.ScalMountPoint` resource types. This property specifies whether to reboot the local system when a resource of these types detects a failure on the local node and the same resource might be able to come online on some other primary or potential primary node.
- A new extension property, `IOTimeout`, is added to the `SUNW.ScalDeviceGroup` resource type. This property defines the length of time at which an I/O probe is considered to have failed, similar to the `IOTimeout` property currently available in the `SUNW.ScalMountPoint` and `SUNW.HASStoragePlus` resource types.

- The `Ping_timeout` property for disk path monitoring accepts higher and lower timeout values.

See the following man pages for details:

- `scdpmd.conf(4)`
- `SUNW.HASStoragePlus(5)`
- `SUNW.ScalDeviceGroup(5)`
- `SUNW.ScalMountPoint(5)`

New clsetup Wizard to Configure the Oracle PeopleSoft Application Server Data Service

The `clsetup` utility now enables you to configure the HA for Oracle PeopleSoft Application Server data services. For details, see [Chapter 2, “Installing and Configuring the HA for PeopleSoft Application Server,”](#) in *Oracle Solaris Cluster Data Service for Oracle PeopleSoft Enterprise Guide*.

New clsetup Wizard to Configure the Oracle WebLogic Server Data Service

The `clsetup` utility now enables you to configure the HA for Oracle WebLogic Server data service. For details, see [Chapter 1, “Installing and Configuring Oracle Solaris Cluster HA for Oracle WebLogic Server,”](#) in *Oracle Solaris Cluster Data Service for Oracle WebLogic Server Guide*.

New Data Service for Oracle External Proxy

A new data service for Oracle External Proxy software is now available. For details, see [Oracle Solaris Cluster Data Service for Oracle External Proxy Guide](#).

New Data Service for SAP NetWeaver

A new data service for SAP NetWeaver software is now available. For details, see [Oracle Solaris Cluster Data Service for SAP NetWeaver Guide](#).

New Data Service for Oracle PeopleSoft Enterprise Process Scheduler

A new data service for the Oracle PeopleSoft Enterprise Process Scheduler is now available for certain batch processing and workload balancing. For details, see [Chapter 3, “Installing and Configuring the HA for PeopleSoft Process Scheduler,”](#) in *Oracle Solaris Cluster Data Service for Oracle PeopleSoft Enterprise Guide*.

New Data Service for Oracle Web Tier

A new data service for Oracle Web Tier software is now available. For details, see *Oracle Solaris Cluster Data Service for Oracle Web Tier Guide*.

Support for Sun ZFS Storage Appliance Data Replication With Geographic Edition

Support is added to use Sun ZFS Storage Appliance data replication in an Oracle Solaris Cluster Geographic Edition configuration. For more information, see *Oracle Solaris Cluster Geographic Edition Remote Replication Guide for Sun ZFS Storage Appliance*.

Restrictions

The following restrictions are in effect as of the time of publication. Contact your Oracle representative to learn whether a feature or product becomes qualified with Oracle Solaris Cluster 3.3 3/13 software at a later date.

- [“Oracle ACFS as a Cluster File System”](#) on page 15
- [“Topology Pages in Oracle Solaris Cluster Manager”](#) on page 16

Oracle ACFS as a Cluster File System

The following are restrictions on the use of Oracle ACFS as a cluster file system in an Oracle Solaris Cluster configuration:

- You can use an Oracle ACFS file system only with Oracle Solaris Cluster HA for Apache, Oracle Solaris Cluster HA for NFS, Oracle Solaris Cluster HA for Oracle, and the database home for Oracle Real Application Clusters (Oracle RAC)..
- You cannot configure an Oracle ACFS file system by using the `clsetup` utility or the configuration wizards. You must configure an Oracle ACFS file system manually by using Oracle Solaris Cluster maintenance commands.

Topology Pages in Oracle Solaris Cluster Manager

Clusters that have more than four nodes could experience a time-out of the topology pages and the pages might fail to render in the Solaris Cluster Manager.

Commands Modified in This Release

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

Compatibility Issues

This section contains the following information about Oracle Solaris Cluster compatibility issues with other products. Contact your Oracle support representative to learn whether a fix becomes available.

- [“Oracle ASM With Solaris Volume Manager Mirrored Logical Volumes”](#) on page 16
- [“ACFS Proxy Resource Goes Offline During Installation of Oracle Binaries \(16581850\)”](#) on page 17
- [“Tavor InfiniBand \(IB\) Transport Cluster Nodes Panic: mutex_enter: bad mutex, lp \(16372022\)”](#) on page 17
- [“Solaris Upgrade to Oracle Solaris 10 8/11 Overwrites the `hostname.adp` Contents \(15882549\)”](#) on page 17
- [“Solaris Volume Manager GUI”](#) on page 18

See also the following information:

- Additional Oracle Solaris Cluster framework compatibility issues are documented in Chapter 1, “Planning the Oracle Solaris Cluster Configuration,” in *Oracle Solaris Cluster Software Installation Guide*.
- Additional Oracle Solaris Cluster upgrade compatibility issues are documented in “Upgrade Requirements and Software Support Guidelines” in *Oracle Solaris Cluster Upgrade Guide*.
- For other known problems or restrictions, see [“Known Issues and Bugs”](#) on page 21.

Oracle ASM With Solaris Volume Manager Mirrored Logical Volumes

Problem Summary: If your Oracle Solaris Cluster HA for Oracle Database or Support for Oracle RAC configuration requires using Oracle ASM with Solaris Volume Manager mirrored logical volumes, you might experience failures of the `SUNW.ScalDeviceGroup` probe. These failures result in a loss of availability of any service that is dependent on the `SUNW.ScalDeviceGroup` resource.

Workaround: You can mitigate the failures by increasing the `IOTimeout` property setting for the `SUNW.ScalDeviceGroup` resource type. See Article 603825.1 at [My Oracle Support](https://support.oracle.com) (<https://support.oracle.com>) for additional information.

ACFS Proxy Resource Goes Offline During Installation of Oracle Binaries (16581850)

Problem Summary: Because of an issue with ACFS resources going offline unexpectedly, the combination of Oracle 12.1 and ACFS with Oracle Solaris Cluster 3.3 3/13 is not currently supported. This issue does not affect Oracle 11.2.0.3 ACFS support.

Workaround: Contact your Oracle support representative to learn whether a fix or workaround is available.

Tavor InfiniBand (IB) Transport Cluster Nodes Panic: mutex_enter: bad mutex, lp (16372022)

Problem Summary: Clusters with Tavor IB Host Channel Adapters (HCAs) panic when attempting to use Reliable Datagram Sockets (RDS) v3 drivers. RDSv3 is not supported on Tavor IB because of hardware limitations. Use the following command to determine if Tavor IB is installed:

```
# /usr/sbin/prtconf -v | grep tavor
```

Workaround: Configure applications on the cluster to use either RDS version 1 or the IP-over-IB protocol. In order to run RDSv3, the Hermon IB HCAs must be installed on the cluster.

Solaris Upgrade to Oracle Solaris 10 8/11 Overwrites the `hostname.adp` Contents (15882549)

Problem Summary: During upgrade from Oracle Solaris 9 to Oracle Solaris version 10 8/11 or 10 1/13, the `/etc/hostname.adp` file is overwritten. This causes networking resources to fail to come online, because the IPMP group configuration was removed during the file overwrite.

Workaround: Perform the following steps for each node you upgrade:

1. Before you upgrade the Oracle Solaris OS, save the contents of the `/etc/hostname.adp` file.
2. After the upgrade but before you reboot the system, copy the saved contents of the pre-upgrade `hostname.adp` files to the current `hostname.adp` files.

3. After you reboot the system, if there was IPMP group information available in the `/etc/hostname.adp` files, verify that IPMP groups are created correctly. Use either the `cluster` command or `ifconfig` commands.

Solaris Volume Manager GUI

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

Accessibility Information

Our goal is to make Oracle products, services, and supporting documentation accessible, with good usability, to the disabled community.

Our documentation includes features that make information available to users of assistive technology. The product documentation is available in HTML format and contains markup to facilitate access by the disabled community. For more information, visit the Oracle Accessibility Program web site at <http://www.oracle.com/us/corporate/accessibility/>.

Supported Products

This section describes the supported software and memory requirements for Oracle Solaris Cluster 3.3 3/13 software as of the time of publication. Contact your Oracle support representative to learn whether additional product versions become supported.

- “Data Services” on page 18
- “File Systems” on page 19
- “Memory Requirements” on page 19
- “Oracle Solaris Operating System” on page 19
- “Oracle VM Server for SPARC” on page 20
- “Availability Suite” on page 20
- “Volume Managers” on page 20

Data Services

Contact your Oracle sales representative for the complete list of supported data services (agents) and application versions.

File Systems

The following table describes the supported file systems for Oracle Solaris 10 on the SPARC and x86 platforms as of the time of publication.

File System	Additional Information
Failover UFS local file system	
ZFS (exported as a data set)	Not supported for the <code>/globaldevices</code> file system
QFS standalone file system and QFS shared file system (only when used to support Oracle Real Application Clusters)	
NFS from supported NAS devices	
VxFS	Requires a minimum of Oracle Solaris 10 1/13 software

Memory Requirements

Oracle Solaris Cluster 3.3 3/13 software requires the following memory requirements for every cluster node:

- Minimum of 1 Gbyte 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.

Oracle Solaris Operating System

Oracle Solaris Cluster 3.3 3/13 software and Quorum Server software requires one of the following versions of the Oracle Solaris OS as of the time of publication:

- **Oracle Solaris 10** – Solaris 10 8/11, Oracle Solaris 10 1/13

Note – Oracle Solaris Cluster 3.3 3/13 software does not support multiple versions of Oracle Solaris software in the same running cluster.

Oracle VM Server for SPARC

This Oracle Solaris Cluster release supports Oracle VM Server for SPARC 2.2 and 3.0 as of the time of publication.

Availability Suite

This Oracle Solaris Cluster release supports Availability Suite 4.0 software as of the time of publication. Support requires a minimum of patch 123246-07 for SPARC or 123247-07 for x86.

Volume Managers

The following volume managers are supported for Oracle Solaris 10 on the SPARC and x86 platforms as of the time of publication:

- Solaris Volume Manager
- Solaris Volume Manager for Sun Cluster
- Veritas Volume Manager (requires a minimum of Oracle Solaris 10 1/13 software)

Features Not Supported in This Release

The following features are not supported in the current Oracle Solaris Cluster release.

Sun Management Center

Sun Management Center is no longer supported in this release. Instead, use Oracle Solaris Cluster Manager to monitor cluster function.

Network Appliance NAS Devices

Network Appliance NAS is not supported in this release. Instead, use Oracle's Sun NAS device or the Sun ZFS Storage Appliance from Oracle.

Socket Direct Protocol (SDP)

Socket Direct Protocol (SDP) over an Oracle Solaris Cluster InfiniBand interconnect is not supported in this release.

Product Localization

The following table describes localization for certain components of Oracle Solaris Cluster 3.3 3/13 software:

Component	Localization
Software Command Line	Japanese, Korean, Simplified Chinese
Software GUI	French, Japanese, Korean, Simplified Chinese, Spanish
Online Help	French, Japanese, Simplified Chinese, Spanish
Man Pages	Japanese

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

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

Known Issues and Bugs

The following known issues and bugs affect the operation of the Oracle Solaris Cluster 3.3 3/13 release. Contact your Oracle support representative to learn whether a fix becomes available. Bugs and issues are grouped into the following categories:

- [“Administration” on page 22](#)
- [“Data Services” on page 26](#)
- [“Geographic Edition” on page 31](#)
- [“Installation” on page 32](#)
- [“Localization” on page 34](#)
- [“Runtime” on page 34](#)

Administration

Cannot Set the Jumbo Frame MTU Size for the clprivnet Interface (16618736)

Problem Summary: The MTU of the cluster `clprivnet` interface is always set to the default value of 1500 and does not match the MTU of the underlying private interconnects. Therefore, you cannot set the jumbo frame MTU size for the `clprivnet` interface.

Workaround: There is no known workaround. Contact your Oracle support representative to learn whether a patch becomes available.

Cluster Check Utility S6708502 Reports RT Process ora_dism as Unsupported (16365116)

Problem Summary: The cluster check utility may report a violation of check S6708502, indicating that real-time process `ora_dism` is not supported for Oracle Solaris Cluster.

Workaround: Ignore the check violation for this specific process. This real-time process is new for Oracle RAC 12c and is allowed for Oracle Solaris Cluster.

Public Net Failure Does Not Fail Over DB Server Resource with SCAN Listener (16231523)

Problem Summary: The HA-Oracle database resource will not fail over when the public network fails when the HA-Oracle database is configured to use the Grid Infrastructure SCAN listener.

Workaround: When using the Oracle Grid Infrastructure SCAN listener with an HA-Oracle database, add a logical host with an IP address that is on the same subnet as the SCAN listener to the HA-Oracle database resource group.

Cluster File System Does Not Support Extended Attributes (15790565)

Problem Summary: Extended attributes are not currently supported by cluster file systems. When a user mounts a cluster file system with the `xattmount` option, the following behavior is seen:

- The extended attribute operations on a regular file fail with an `ENOENT` error.
- The extended attribute operations on a directory results in normal operations on the directory itself.

Therefore, any program accessing the extended attributes of files in a cluster file system might not get the expected results.

Workaround: Mount a cluster file system with the `noxttmount` option.

Resource Group Does Not Fail Over When Failover_mode Is Set to SOFT During a Public Interface Failure (15711034)

Problem Summary: If a failover data service, such as HA for Oracle, is configured with the ScalMountpoint resource to probe and detect NAS storage access failure, and the network interface is lost, such as due to a loss of cable connection, the monitor probe hangs. If the Failover_mode property of the data service resource is set to SOFT, this results in a stop-failed status and the resource does not fail over. The associated error message is similar to the following:

```
SC[SUNW.ScalMountPoint:3,scalmnt-rg,scal-oradata-11g-rs,/usr/cluster/lib/rgm
  /rt/scal_mountpoint/scal_mountpoint_probe]:
Probing thread for mountpoint /oradata/11g is hanging for timeout period
  300 seconds
```

Workaround: Change the Failover_mode property on the data service resource to HARD.

```
# clresource set -p Failover_mode=HARD ora-server-rs
# clresource show -v ora-server-rs | grep Failover_mode
  Failover_mode: HARD
```

Unable to Register Resource Type SUNW.scalable_acfs_proxy in a Zone Cluster (15700050)

Problem Summary: The current implementation requires an RTR file, rather than a symbolic link to the file, to be present in /usr/cluster/lib/rgm/rtreg.

Workaround: Perform the following commands as superuser on one node of the global cluster.

```
# cp /opt/SUNWscor/oracle_asm/etc/SUNW.scalable_acfs_proxy /usr/cluster/lib/rgm/rtreg/
# clrt register -Z zoneclustername SUNW.scalable_acfs_proxy
# rm /usr/cluster/lib/rgm/rtreg/SUNW.scalable_acfs_proxy
```

The clzonecluster boot, reboot, and halt Subcommands Fail if Any One of the Cluster Nodes Is Not in the Cluster (15812235)

Problem Summary: The clzonecluster boot, reboot, and halt subcommands fail, even if one of the cluster nodes is not in the cluster. An error similar to the following is displayed:

```
root@pnode1:~# clzc reboot zoneclustername
clzc: (C827595) "pnode2" is not in cluster mode.
clzc: (C493113) No such object.

root@pnode1:~# clzc halt zoneclustername
clzc: (C827595) "pnode2" is not in cluster mode.
clzc: (C493113) No such object.
```

The clzonecluster boot, reboot, and halt subcommands should skip over nodes that are in noncluster mode, rather than fail.

Workaround: Use the following option with the `clzonecluster boot` or `clzonecluster halt` commands to specify the list of nodes for the subcommand:

```
-n nodename[, ...]
```

The `-n` option allows running the subcommands on the specified subset of nodes. For example, if, in a three-node cluster with the nodes `pnode1`, `pnode2`, and `pnode3`, the node `pnode2` is down, you could run the following `clzonecluster` subcommands to exclude the down node:

```
clzonecluster halt -n pnode1,pnode3 zoneclustername
clzonecluster boot -n pnode1,pnode3 zoneclustername
clzonecluster reboot -n pnode1,pnode3 zoneclustername
```

Using `chmod` to Set `setuid` Permission Returns Error in a Non-Global Zone on PxFs Secondary Server (15697690)

Problem Summary: The `chmod` command might fail to change `setuid` permissions on a file in a cluster file system. If the `chmod` command is run on a non-global zone and the non-global zone is not on the PxFs primary server, the `chmod` command fails to change the `setuid` permission.

For example:

```
# chmod 4755 /global/oracle/test-file
chmod: WARNING: can't change /global/oracle/test-file
```

Workaround: Do one of the following:

- Perform the operation on any global-cluster node that accesses the cluster file system.
- Perform the operation on any non-global zone that runs on the PxFs primary node that has a loopback mount to the cluster file system.
- Switch the PxFs primary to the global-cluster node where the non-global zone that encountered the error is running.

Cannot Create a Resource From a Configuration File With Non-Tunable Extension Properties (15658184)

Problem Summary: When you use an XML configuration file to create resources, if any of the resources have extension properties that are not tunable, that is, the `Tunable` resource property attribute is set to `None`, the command fails to create the resource.

Workaround: Edit the XML configuration file to remove the non-tunable extension properties from the resource.

Disabling Device Fencing While Cluster Is Under Load Results in Reservation Conflict (15608862)

Problem Summary: Turning off fencing for a shared device with an active I/O load might result in a reservation conflict panic for one of the nodes that is connected to the device.

Workaround: Quiesce I/O to a device before you turn off fencing for that device.

Autodiscovery Does Not Work on LDoms With Hybrid I/O (15581535)

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

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

EMC SRDF and Hitachi TrueCopy Should Reject Switchover When Replicated Device-Group Status Will Cause Switchover and Switchback to Fail (15538295)

Problem Summary: If a Hitachi TrueCopy device group whose replica pair is in the COPY state, or an EMC SRDF device group whose replica pair is split, attempts to switch the device group over to another node, the switchover fails. Furthermore, the device group is unable to come back online on the original node until the replica pair is been returned to a paired state.

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

Removing Nodes from the Cluster Configuration Can Result in Node Panics (15500647)

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

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

More Validation Checks Needed When Combining DIDs (15423531)

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

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

Data Services

sun.storage_proxy.type Resource in Offline State After uadmin 5 1 (16910145)

Problem Summary: During an ungraceful shutdown of a cluster node, such as a node panic, the Oracle Clusterware `sun.storage-proxy-resource` of type `sun.storage_proxy.type` might stay offline upon node bootup. This in turn will cause the Oracle Solaris Cluster RAC server proxy resource to stay offline.

Workaround: Perform the following steps:

1. **Bring up the ACFS storage proxy resource manually.**

```
# crsctl stop res sun.storage-proxy-resource -n nodename  
# crsctl start res sun.storage-proxy-resource -n nodename
```

2. **Bring online the Oracle Solaris Cluster RAC server-proxy resource.**

```
# clresourcegroup online rac-server-proxy-resource-group
```

Active-Standby Configuration Not Supported for HA for TimesTen (16861602)

Problem Summary: The TimesTen active-standby configuration requires an integration of Oracle Solaris Cluster methods in the TimesTen `ttCwadmin` utility. This integration has not yet occurred, even though it is described in the [Oracle Solaris Cluster Data Service for Oracle TimesTen Guide](#). Therefore, do not use the TimesTen active-standby configuration with Oracle Solaris Cluster HA for TimesTen and do not use the TimesTen `ttCwadmin` utility on Oracle Solaris Cluster.

The Oracle Solaris Cluster TimesTen data service comes with a set of resource types. Most of these resource types are meant to be used with TimesTen active-standby configurations. You must use only the `ORCL.TimesTen_server` resource type for your highly available TimesTen configurations with Oracle Solaris Cluster.

Workaround: Do not use the TimesTen active-standby configuration.

The `clzonecluster halt zonecluster` Command Followed by `clzonecluster boot zonecluster` Causes Node Boot Failure in a Zone Cluster (16398199)

Problem Summary: If you run the `clzonecluster halt zonecluster` command, followed by the `clzonecluster boot zonecluster` command, one or more of the nodes fail to boot with the following error:

```
root@node1:/# clzonecluster boot zc1
Waiting for zone boot commands to complete on all the nodes of the
zone cluster "zc1"...
clzc: (C215301) Command execution failed on node node2.
zoneadm: zone 'zc1': These file-systems are mounted on subdirectories
of /gpool/zones/zone1/root:
zoneadm: zone 'zc1': /gpool/zones/zone1/root/u01
zoneadm: zone 'zc1': call to zoneadmd failed
```

The zone cluster node will not boot and the `clzonecluster status` command shows the node(s) as offline.

Workaround: Unmount the file system in the global zone of the offline node(s): `/usr/sbin/umount/gpool/zones/zone1/root/u01` (in the case above) and run the following command in the global zone of any node of the zone cluster:
`/usr/cluster/bin/clzonecluster boot -n offline-node zonecluster`. Verify that the offline node(s) are now online by running the `/usr/cluster/bin/clzonecluster status` command.

Disabled LDom Resource Keeps the LDom in a Bound State (16323723)

Problem Summary: When the HA for Oracle VM Server for SPARC (HA for Logical Domains) resource is disabled during manual maintenance operations, the zpool export fails. This failure occurs because the bound state of the logical domain keeps the ZFS zpool, which is dependent on the failover ZFS resource, in a busy state. Switchovers and failovers are not affected.

Workaround: Perform the following steps:

1. Release resources from the logical domain.

```
# ldm unbind-dom ldom
```

2. Clear the HASP resource that is in a STOP_FAILED state.

For more information, see [“Clearing the STOP_FAILED Error Flag on Resources”](#) in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

Manual Umount of ZFS File System Does Not Get Remounted Automatically (16319552)

Problem Summary: The HAStoragePlus probe does not automatically remount a Solaris ZFS filesystem, if it has been unmounted.

Workaround: Provide a mount point to the dataset, and then any manually unmounted file systems are automatically remounted by the HAStoragePlus probe. For example:

```
# zfs list
NAME                USED  AVAIL  REFER  MOUNTPOINT
pool-1              414K  417G   32K    none
pool-1/test1       31.5K  417G  31.5K  /testmount1
pool-1/test2       31.5K  417G  31.5K  /testmount2
```

If pool-1 is given a mount point, then any manually unmounted file systems are automatically remounted by the HAStoragePlus probe.

```
# zfs set -p mountpoint=/pool-1 pool-1
# zfs list
NAME                USED  AVAIL  REFER  MOUNTPOINT
pool-1              414K  417G   32K    /pool-1
pool-1/test1       31.5K  417G  31.5K  /testmount1
pool-1/test2       31.5K  417G  31.5K  /testmount2
```

Oracle HTTP Server (OHS) Agent Calls /bin/grep -q, Which Is an Illegal Option for This grep Version (16270576)

Problem Summary: When you stop the HA-OHS data service, the */bin/grep: illegal option -q* message appears even though the resource goes offline successfully.

Workaround: Ensure that the OHS processes are not running before bringing the OHS application under cluster control. If the processes are not running outside the cluster control, you can ignore this message.

Errors Occur When Deleting an HA LDom Resource (16229698)

Problem Summary: Error messages related to removal of a CCR entry (domain configuration) appear when the resource is being deleted.

Workaround: These error messages are harmless and can be ignored.

The clsetup WebLogic Server Wizard Does Not Create a Resource for a Failover File System in a Zone Cluster Zone (15943191)

Problem Summary: If the WebLogic Server application is configured on a failover file system, the wizard fails to create a resource. If the WebLogic Server instances are configured to listen on "All IP Addresses", the wizard also fails to create a resource.

Workaround: If the WebLogic Server application is configured on a failover filesystem and the wizard fails to create a resource, manually create the HA-WLS resource for a failover file system. If the WebLogic Server instances are configured to listen on "All IP Addresses" and the wizard fails to create a resource, configure the instances to listen only on logical hosts as recommended in the documentation or manually create the HA-WLS resource.

Wizard Unable to Create Zone Cluster rac-framework Resource Group if the Zone Cluster Has Fewer Nodes Than Its Global Zone Nodes (15825830)

Problem Summary: Configuring RAC framework using the data service wizard fails in a zone cluster if the wizard is run from a node that is not a part of the zone cluster.

Workaround: If configuring the data service in a zone cluster, run the wizard from one of the cluster nodes that is hosting the zone cluster.

Oracle Solaris 10 Samba Patch 119757-20 or 119758-20 Impacts the HA-Samba Data Service (15757158)

Problem Summary: After applying the Samba patch 119757-20 (SPARC) or 119758-20 (x86), the binary locations change from `/usr/sfw/bin` to `/usr/bin` and `/usr/sfw/lib` to `/usr/lib/samba`. This breaks the Oracle Solaris Cluster Data Service for Samba.

Workaround: If the patches listed above have been installed, then the Oracle Solaris Cluster Data Service for Samba resource needs to be reregistered (you must remove it and register it again). The `/opt/SUNWscsmb/util/samba_config` file must specify the new binary locations as described above. After the `samba_config` file has been changed, the `/opt/SUNWscsmb/util/samba_register` file must be executed to register the resource again.

When set Debug_level=1, pas-rg Fails Over to Node 2 And Cannot Start on Node 1 Anymore (15803844)

Problem Summary: If you set the `Debug_level` property to 1, a start of a dialogue instance resource is impossible on any node.

Workaround: Use `Debug_level=2`, which is a superset of `Debug_level=1`.

HASStoragePlus Resource Configured in Scalable Resource Group with Cluster File System Stays at Starting State Indefinitely (15649193)

Problem Summary: When the `/etc/vfstab` file entry for a cluster file system has a `mount-at-boot` value of `noand` and the cluster file system is configured in a `SUNW.HASStoragePlus` resource that belongs to a scalable resource group, the `SUNW.HASStoragePlus` resource fails to come online. The resource stays in the `Starting` state until `prenet_start_method` is timed out.

Workaround: In the `/etc/vfstab` file's entry for the cluster file system, set the `mount-at-boot` value to `yes`.

Scalable Applications Are Not Isolated Between Zone Clusters (15611122)

Problem Summary: If scalable applications configured to run in different zone clusters bind to `INADDR_ANY` and use the same port, then scalable services cannot distinguish between the instances of these applications that run in different zone clusters.

Workaround: Do not configure the scalable applications to bind to `INADDR_ANY` as the local IP address, or bind them to a port that does not conflict with another scalable application.

Running `clnas add` or `clnas remove` Commands on Multiple Nodes at the Same Time Can Cause a Problem (15533979)

When adding or removing a NAS device, running the `clnas add` or `clnas remove` command on multiple nodes at the same time might corrupt the NAS configuration file.

Workaround: Run the `clnas add` or `clnas remove` command on one node at a time.

The `clresourcegroup add-node` Command Triggers an `HAStoragePlus` Resource to a Faulted State (15391480)

Problem Summary: When a native brand non-global zone is added to the node list of a resource group that contains an `HAStoragePlus` resource with ZFS pools configured, the `HAStoragePlus` resource might enter the `Faulted` state. This problem happens only when the physical node that hosts the native zone is part of the resource-group node list.

Workaround: Restart the resource group that contains the faulted `HAStoragePlus` resource.

```
# clresourcegroup restart faulted-resourcegroup
```

RAC Wizard Failing With "ERROR: Oracle ASM is either not installed or the installation is invalid!" (15814020)

Problem Summary: The Oracle RAC configuration wizard fails with the message, `ERROR: Oracle ASM is either not installed or the installation is invalid!`.

Workaround: Ensure that the "ASM" entry is first within the `/var/opt/oracle/oratab` file, as follows:

```
root@phys-schost-1:~# more /var/opt/oracle/oratab
...
+ASM1:/u01/app/11.2.0/grid:N          # line added by Agent
MOON:/oracle/ora_base/home:N
```

Geographic Edition

Geographic Edition Hitachi TrueCopy Protection Group Takeover Fails When the Primary Cluster Loses Access to Storage (16013386)

Problem Summary: The takeover operation fails when the primary cluster loses access to the storage device.

Workaround: Bring the primary cluster with lost access to storage down.

Geographic Edition Disables ASM Resources in Incorrect Order, Resulting in Error Messages (16004984)

Problem Summary: When Geographic Edition switches a protection group (PG) to have the secondary role, it unmanages ASM device groups incorrectly. If the cluster is then restarted and if the LUNs are read-write, these devices groups are incorrectly re-enabled. When Geographic Edition is restarted, writes to these LUNs will be disabled and the user might see several fatal write errors on the system console. These errors do not indicate a serious problem and can be ignored. Geographic Edition will operate correctly.

Workaround: Ignore the messages.

Cluster Node Does Not Have Access to ZFS Storage Appliance Projects or iSCSI LUNs (15924240)

Problem Summary: If a node leaves the cluster when the site is the primary, the projects or iSCSI LUNs are fenced off. However, after a switchover or takeover when the node joins the new secondary, the projects or iSCSI LUNs are not unfenced and the applications on this node are not able to access the file system after it is promoted to the primary.

Workaround: Reboot the node.

Multiple Notification Emails Sent from Global Cluster When Zone Clusters Are in Use (15746425)

Problem Summary: If Oracle Solaris Cluster Geographic Edition is configured in a zone cluster, duplicate notification emails about loss of connection to partner clusters are sent from both the zone cluster and the global cluster. Emails should only be sent from the zone cluster.

Workaround: This is a side effect of the cluster event handling. It is harmless and you should ignore the duplicate emails.

DR State Stays Reporting unknown on One Partner (15808054)

Problem Summary: DR state stays reporting unknown, although DR resources are correctly reporting replication state.

Workaround: Run the `geopg validate protection-group` command to force a resource-group state notification to the protection group.

ZFS Storage Appliance Protection Group Creation And Validation Fail if Project Replication Is Stopped by Using the BUI (15797609)

Problem Summary: If you use the browser user interface (BUI) to stop replication, the protection group goes to a configuration error state when protection-group validation fails.

Workaround: From the BUI, perform the following actions to stop replication:

1. Under the Shares tab, select the project being replicated.
2. Click on the Replication tab and select the Scheduled option.
3. Wait until the status changes to manual, then click the Enable/Disable button.

Installation

Centralized Install with the `scinstall` Utility Fails to Configure the Cluster When DES Authentication is Enabled and the Node Names Are Specified as Fully Qualified Host Names (16228318)

Problem Summary: When you use the centralized install, the `scinstall` utility fails to configure the cluster when DES authentication is enabled and nodes are specified as fully qualified host names. An error message similar to the following appears:

```
Updating file ("ntp.conf.cluster") on node <FQ-host-name> ... failed
scinstall: Failed to configure ("ntp.conf.cluster") on node <FQ-host-name>
scinstall: scinstall did NOT complete successfully!
```

Workaround: Rerun the `scinstall` utility and this time select the option to configure one node at a time. Specify the node name without the domain name. If you are configuring a two-node cluster, the quorum configuration will fail causing the install mode to not be reset. Manually reset the install mode in that case, after the nodes boot up into cluster mode.

The Zone Cluster Wizard Does Not Provide a Method to Enter Data on Nodes Using Oracle Solaris 10 OS (15874818)

Problem Summary: A zone cluster can be in a Ready-Offline state when it is booted because of incomplete Oracle Solaris system configuration in the zones of the zone cluster. The zones of the zone cluster are in the interactive system configuration mode to wait for the input. This occurs when there is no system configuration file (`/etc/sysidcfg`) or the file does not contain all required system configuration properties in the global zone on the cluster node.

Workaround: Before installing a zone cluster, create the `/etc/sysidcfg` file and specify all required system configuration properties on all cluster nodes. The configuration properties in the file are used to do the Oracle Solaris system configuration automatically during the first boot

of the zone cluster. The list of required Oracle Solaris system configuration properties can vary according to the Oracle Solaris OS version. See the [Oracle Solaris Administration: Basic Administration](#) for more details. Alternatively, after the zone cluster is installed and booted, use `zlogin -C zone` to log in on all nodes of the zone cluster to manually complete the Oracle Solaris system configuration.

Running the `scinstall -u` Command Leaves 'installed' Zones in 'mounted' State (15817301)

Problem Summary: Running the `scinstall -u` command leaves 'installed' zones in a 'mounted' state. This state causes an issue for live upgrade when the system is rebooted as it fails to fix the zonepaths for the alternate boot environment.

Workaround: Perform the following steps:

1. Run the `svcadm disable zones` command.
All running zones should now be in a mounted state. The service might end up in maintenance after 100 seconds, but this is not a problem.
2. Run the `zoneadm -z zonename unmount` command for all zones.
3. Type `init 6`.

The Installer Deletes the Existing Package Corresponding to Ops Center Agent JavaDB Database (15646335)

Problem Summary: The Oracle Enterprise Manager Ops Center Agent for Oracle Solaris 10 uses JavaDB software for its configuration database. When installing the Oracle Solaris Cluster software by using the `installer` utility, the JavaDB software package is reinstalled, causing an existing agent configuration database to be deleted.

The following error messages are reported from the Ops Center Agent as a result of the package getting removed:

```
java.sql.SQLException: Database '/var/opt/sun/xvm/agentdb' not found.
at org.apache.derby.impl.jdbc.SQLExceptionFactory40.getSQLException(Unknown Source)
at org.apache.derby.impl.jdbc.Util.newEmbedSQLException(Unknown Source)
at org.apache.derby.impl.jdbc.Util.newEmbedSQLException(Unknown Source)
```

The Agent is broken now and needs to be unconfigured or configured.

Workaround: Manually install on all cluster nodes the following additional JavaDB packages from the Oracle Solaris Cluster media:

- `SUNWjavadb-demo`
- `SUNWjavadb-javadoc`
- `SUNWjavadb-docs`
- `SUNWjavadb-client`

Running the `installer` utility does not remove the existing JavaDB database packages.

Localization

Result of System Requirements Checking Is Wrong (12185261)

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

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

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

Runtime

The `clzonecluster cancel` Command Can Cause an Error (15950187)

Problem Summary: The `clzonecluster` interactive configuration (opened by the `clzonecluster configure` command `zcname`) can crash in some circumstances when the `cancel` subcommand is issued. The *Error executing zone configure command* error message is displayed.

Workaround: You can safely ignore this problem. Only unsaved configuration data is lost due to the problem. To avoid a crash of the configuration utility, do not use the `cancel` command.

`sc_delegated_restarter` Does Not Take Into Account Environment Variable Set in Manifest (15795184)

Problem Summary: Any environment variables that are specified in the service manifest are not recognized when the service is put under `SUNW.Proxy_SMF_failover` resource type control.

Workaround: Modify the service methods to set the environment variables directly.

Unable to Re-enable Transport Interface After Disabling With `ifconfig unplumb interface` (15770198)

Problem Summary: Cluster transport paths go offline with accidental use of the `ifconfig unplumb` command on the private transport interface.

Workaround: Disable and re-enable the cable that the disabled interface is connected to.

1. Determine the cable to which the interface is connected.

- ```
/usr/cluster/bin/clinterconnect show | grep Cable
```
2. Disable the cable for this interface on this node.

```
/usr/cluster/bin/clinterconnect disable cable
```
  3. Re-enable the cable to bring the path online.

```
/usr/cluster/bin/clinterconnect enable cable
```

## Failure of Logical Hostname to Fail Over Caused by getnetmaskbyaddr() (15733435)

**Problem Summary:** Logical hostname failover requires getting the netmask from the network if nis/ldap is enabled for the netmasks name service. This call to `getnetmaskbyaddr()` hangs for a while due to CR 7051511, which might hang long enough for the Resource Group Manager (RGM) to put the resource in the FAILED state. This occurs even though the correct netmask entries are in the `/etc/netmasks` local files. This issue affects only multi-homed clusters, such as cluster nodes that reside on multiple subnets.

**Workaround:** Configure the `/etc/nsswitch.conf` file, which is handled by an SMF service, to only use files for netmasks lookups.

```
/usr/sbin/svccfg -s svc:/system/name-service/switch setprop config/netmask = astring:"files"
/usr/sbin/svccadm refresh svc:/system/name-service/switch
```

## Patches and Required Firmware Levels

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

- “[Applying an Oracle Solaris Cluster 3.3 3/13 Core Patch](#)” on page 36
- “[Removing an Oracle Solaris Cluster 3.3 3/13 Core Patch](#)” on page 36
- “[Patch Management Tools](#)” on page 37
- “[Patch for Cluster Support for Sun StorageTek 2530 Array](#)” on page 37
- “[My Oracle Support](#)” on page 38

---

**Note** – If you are upgrading to Oracle Solaris Cluster 3.3 3/13 software, see [Oracle Solaris Cluster Upgrade Guide](#). Applying a Oracle Solaris Cluster 3.3 3/13 Core patch does *not* provide the same result as upgrading the software to the Oracle Solaris Cluster 3.3 3/13 release.

---



---

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

---

You must be a registered My Oracle Support user to view and download the required patches for the Oracle Solaris Cluster product. If you do not have a My Oracle Support account, contact your Oracle service representative or sales engineer, or register online at <http://support.oracle.com>.

## Applying an Oracle Solaris Cluster 3.3 3/13 Core Patch

Complete the following procedure to apply the Oracle Solaris Cluster 3.3 3/13 core patch. Ensure that all nodes of the cluster are maintained at the same patch level.



**Caution** – If an Oracle Solaris Cluster 3.3 3/13 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 Oracle Solaris Cluster 3.3 3/13 core patch permanently to your cluster.

### ▼ How to Apply the Oracle Solaris Cluster 3.3 3/13 Core Patch

- 1 Install the patch using the usual rebooting patch procedure for a core patch.
- 2 Verify that the patch has been installed correctly on all nodes and is functioning properly.
- 3 Perform Resource Type upgrades to any new versions of the resource type available.

Run the `clsetup` command to obtain the list of new resource types. For more details, see “Upgrading a Resource Type” in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

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

## Removing an Oracle Solaris Cluster 3.3 3/13 Core Patch

Complete the following procedure to remove the Oracle Solaris Cluster 3.3 3/13 core patch.

### ▼ How to Remove an Oracle Solaris Cluster 3.3 3/13 Core Patch

- 1 List the resource types on the cluster.  

```
clresourcetype list
```
- 2 If you have upgraded to any new resource types after applying the core patch, follow the directions in “How to Remove a Resource Type” in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

- 3 **Remove the Oracle Solaris Cluster 3.3 3/13 core patch from each node on which you installed the patch.**

```
patchrm patch-id
```

- 4 **Reboot into cluster mode all of the nodes from which you removed the Oracle Solaris Cluster 3.3 3/13 core patch.**

Rebooting all of the nodes from which you removed the Oracle Solaris Cluster 3.3 3/13 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.

- 5 **Reboot any remaining nodes into cluster mode.**

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

## Patch Management Tools

Information about patch management options for the Oracle Solaris OS is available at the web site for Oracle Enterprise Manager Ops Center (formerly Sun Ops Center) at <http://www.oracle.com/technetwork/oem/ops-center>.

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

- Information for using the Oracle Solaris patch management utility, `patchadd`, is provided in Chapter 23, “Managing Patches,” in *Oracle Solaris Administration: Basic Administration*.
- Information for using Oracle Solaris Live Upgrade to apply patches is provided in the Solaris installation guide for Live Upgrade and upgrade planning at *Oracle Solaris 10 1/13 Installation Guide: Live Upgrade and Upgrade Planning*.

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 *Oracle Solaris Cluster System Administration Guide* 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.

## Patch for Cluster Support for Sun StorageTek 2530 Array

The Sun StorageTek Common Array Manager (CAM) software, minimum required Version 6.0.1, provides SCSI3 or PGR support for the Sun StorageTek 2530 array for up to three nodes.

The patch is not a required upgrade for the Sun StorEdge 6130, 2540, 6140, and 6540, and StorageTek FLX240, FLX280 and FLX380 platforms. The CAM patch is available from My Oracle Support.

## My Oracle Support

The My Oracle Support 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 site at <http://support.oracle.com> for the most current matrixes of supported software, firmware, and patch revisions.

Before you install Oracle Solaris Cluster 3.3 3/13 software and apply patches to a cluster component (Oracle Solaris OS, Oracle Solaris 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 11, “Patching Oracle Solaris Cluster Software and Firmware,”](#) in *Oracle Solaris Cluster System Administration Guide*.

## Oracle Solaris Cluster 3.3 3/13 Documentation Set

The Oracle Solaris Cluster 3.3 3/13 user documentation is available in PDF and HTML format at the following web site:

[http://docs.oracle.com/docs/cd/E37745\\_01/index.html](http://docs.oracle.com/docs/cd/E37745_01/index.html)

## Documentation Addendum

This section discusses errors or omissions for documentation, online help, or man pages in the Oracle Solaris Cluster 3.3 3/13 release.

- “Software Installation Guide” on page 38
- “HA for Oracle Guide” on page 39
- “HA for Oracle RAC Guide” on page 39
- “HA for Oracle Web Tier Guide” on page 42
- “Man Pages” on page 43

## Software Installation Guide

The *Oracle Solaris Cluster Software Installation Guide* contains the following errors:

- **“Global Fencing” in Chapter 1** – The value `prefer3` is mistakenly described as the default global-fencing setting. The correct default value is `pathcount`.
- **“How to Configure an HAStoragePlus Resource for a Cluster File System that is Used by Non-Global Zones”** – In Step 6, an incorrect property is used. Instead of the `Network_resources_used` property, use the `Resource_dependencies_offline_restart` property to specify the HAStoragePlus resource. This property dependency makes the application restart when the HAStoragePlus resource with monitoring enabled is restarted.

## HA for Oracle Guide

The *Oracle Solaris Cluster Data Service for Oracle Guide* contains the following omission:

**Using Oracle ASM in both the global cluster and a zone cluster** – If you use Oracle ASM in a global cluster as well as in a zone cluster configured in that global cluster, you must ensure that, in each particular cluster, Oracle ASM can see only those devices that are intended for its use, whether in the global zone or in a zone cluster. If Oracle ASM can see devices that are used by Oracle ASM in a different cluster, this can cause start problems for Oracle Clusterware or Grid Infrastructure, because Oracle ASM sees the devices as already mounted elsewhere.

## HA for Oracle RAC Guide

The *Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide* contains the following errors and omissions:

- **Using Oracle ASM in both the global cluster and a zone cluster** – If you use Oracle ASM in a global cluster as well as in a zone cluster configured in that global cluster, you must ensure that, in each particular cluster, Oracle ASM can see only those devices that are intended for its use, whether in the global zone or in a zone cluster. If Oracle ASM can see devices that are used by Oracle ASM in a different cluster, this can cause start problems for Oracle Clusterware or Grid Infrastructure, because Oracle ASM sees the devices as already mounted elsewhere.
- **“Storage Management Schemes for Oracle Files”** – In Table 1-2, the cell for Oracle Clusterware binary files with an Oracle ACFS file system should show a minus (-) symbol instead of a plus (+) symbol. An ACFS file system cannot store Oracle Clusterware binary files.
- **“Types of Oracle Files That You Can Store on an Oracle ACFS File System”** – Oracle Clusterware binary files is incorrectly included in the list of file types.
- **“How to Create Oracle Solaris Cluster Resources in the Global Cluster for Interoperation With Oracle 10g or 11g”** – In Step 4, the `rg_affinities` property is incorrect. The optional `db-storage-rg` resource group should be specified as a strong positive affinity, as follows:

```
-p rg_affinities=++rac-fmwk-rg[, ++db-storage-rg]
```

- **“How to Create an Oracle ASM Resource for Oracle Clusterware Interoperation With Oracle Solaris Cluster Software”** – The command in Step 8b is incorrect. The correct command is as follows:

```
Grid_home/bin/crs_register sun.node.asm-dg-rs -update -r crs-asm-instance
```

- **“How to Create an Oracle Grid Infrastructure Resource for Interoperation With Oracle Solaris Cluster”** – In Step 3, the description of the HOSTING\_MEMBERS property should be, “The list of nodes in the cluster membership, separated by spaces”.

For Steps 4 through 8, the correct instructions are as follows:

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

```
Grid_home/bin/osdbagrp
```

**5. Set the primary group of the Oracle Grid Infrastructure storage\_proxy resource to the group determined in Step 4.**

```
Grid_home/bin/crsctl setperm resource sun.sc-resource -g "griddba"
```

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

```
oracle_home/bin/osdbagrp
```

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

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

```
Grid_home/bin/crsctl setperm resource sun.sc-resource -u "group:dba-group:r-x"
```

**8. Bring online the Oracle Grid Infrastructure storage\_proxy resource.**

```
Grid_home/bin/crsctl start resource sun.sc-resource
```

- **Upgrading Oracle Grid Infrastructure or Oracle Database software** – If you need to upgrade Oracle Grid Infrastructure or Oracle Database software but are not also upgrading Oracle Solaris or Oracle Solaris Cluster software, you must first remove the grid or database software from Oracle Solaris Cluster control. This includes any zone cluster resources that depend on the resources in the global cluster.

The following instructions describe the steps to perform before and after you upgrade the grid or database software.

- 1. From one node, identify the Oracle Grid Infrastructure and Oracle Database resources that you need to disable.**

```
clresource -t resource-type list
```

The following are the resource types that might be used:

- SUNW.asm\_diskgroup
- SUNW.oracle\_listener
- SUNW.oracle\_rac\_server
- SUNW.oracle\_server

- SUNW.scalable\_acfs\_proxy
  - SUNW.scalable\_asm\_diskgroup\_proxy
  - SUNW.scalable\_asm\_instance
  - SUNW.scalable\_asm\_instance\_proxy
  - SUNW.scalable\_rac\_listener
  - SUNW.scalable\_rac\_server
  - SUNW.scalable\_rac\_server\_proxy
2. **From one node, use the `clresource` command to remove any dependencies by other resources on grid or database resources or both.**

```
clresource set -p Resource_dependencies==dependent-resource application-resource
```

*Resource\_dependencies*

Specifies a resource on which this resource has a dependency.

`-=`

Deletes the specified value.

*dependent-resource*

Specifies the grid or database resource.

*application-resource*

Specifies the application resource that depends on the grid or database resource.

3. **On one node at a time, remove Oracle Solaris Cluster control of resources associated with the Grid Infrastructure or Oracle Database software you intend to upgrade.**

Perform these steps on each node that runs Oracle Grid Infrastructure or Oracle Database software, or both.

- a. **Disable the resources.**

```
clresource disable -n node resource
```

`-n node`

Specifies the node where the resource is to be disabled.

*resource*

Specifies the grid or database resource to disable.

---

**Note** – Do not disable the RAC framework resource, which uses the `SUNW.rac_framework` resource type. This resource is necessary for the operation of Oracle RAC.

---

- b. **After all associated resources on the node are disabled, manually restart the Grid Infrastructure or Oracle Database software, or both, on the nodes where the software was disabled.**

Use the standard Oracle grid or database commands. This action ensures that the components remain outside Oracle Solaris Cluster control when you upgrade the grid or database software.

4. **From one node, upgrade the Oracle Grid Infrastructure or Oracle Database software, or both.**

Follow procedures provided in the Oracle Grid Infrastructure or Oracle Database documentation.

5. **If the upgraded software requires an addition or change of any extension properties, from one node update the affected grid or database resources, or both.**

```
clresource set -p property=property-value {-p property=property-value ...} \
resource
```

-p *property=value* Specifies the setting of a given property.

*resource* Specifies the resource that contains the affected property.

6. **From one node, enable the Oracle Solaris Cluster resources that are associated with Oracle Grid Infrastructure or Oracle Database software, or both.**

```
clresource enable resource
```

Oracle Solaris Cluster control of grid or database resources is restored.

7. **From one node, restore the resource dependencies on the grid or database resources.**

```
clresource set -p Resource_dependencies+=dependent-resource application-resource
```

*Resource\_dependencies*

Specifies a resource on which the application resource has a dependency.

+=

Adds the specified value.

*dependent-resource*

Specifies the grid or database resource.

*application-resource*

Specifies the application resource that depends on the grid or database resource.

## HA for Oracle Web Tier Guide

The *Oracle Solaris Cluster Data Service for Oracle Web Tier Guide* contains the following errors:

- The following note is incorrect and should be ignored:

---

**Note** – You need to install the HA for Oracle Web Tier packages in the global cluster and not in the zone cluster.

---

It is supported to install HA for Oracle Web Tier packages directly in a zone cluster.

- Because this data service is introduced in the Oracle Solaris Cluster 3.3 3/13 release, there is no previous resource type version to upgrade from. Therefore, information in the section “Information for Registering the New Resource Type Version” is extraneous and can be ignored.

## Man Pages

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

### Multiple Man Pages

Some man pages, such as `clquorum(1CL)`, contain information about Network Appliance NAS devices. This information is extraneous and should be ignored.

