

Oracle® Solaris Cluster 4.2 Release Notes



Part No: E39651
September 2014, E39651-02

Copyright © 2000, 2014, 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.

Copyright © 2000, 2014, Oracle et/ou ses affiliés. Tous droits réservés.

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

Using This Documentation	11
Oracle® Solaris Cluster 4.2 Release Notes	13
What's New in the Software	13
Former Features Not Included in the Oracle Solaris Cluster 4.2 Software	17
Product Localization Notes	17
Configuring Accessibility Support for Oracle Solaris Cluster Manager	18
<accessibility-mode> Element	18
<accessibility-profile> Element	18
▼ How to Configure Accessibility Support in the trinidad-config.xml File	19
Commands Modified in This Release	21
Compatibility Notes	23
Compatibility Issues	23
Restriction of Support for Solaris Volume Manager Multi-Owner Disk Sets	23
Creation of labeled Brand of Zone Clusters Fail (19074503)	23
show configuration Attempts Pile Up and Lead to ORA-00020 (18746643)	24
Cacao Creates a Huge File and Host Runs Out of Root Space Quickly (18099260)	24
The Node Panics With FreeLnode Message While Starting the Oracle Solaris Clusterware Inside the solaris10 Brand of Zone Cluster (19032049)	24
VNIC Names Longer Than 16 Characters Cause Problems (17362337)	25
sdplib: ibt_get_ip_paths With APM Failed Status 25 on Oracle Solaris Cluster When Port Is Down (18692783)	25
Zone Cluster Does Not Recognize *dsk WildCard With svm /dev/md (18167722)	25
osysmond Core Dumps in Solaris10 Branded Zone During GI root.sh and Starting of CRS (14456069)	26

Oracle Clusterware Fails to Create All SIDs for ora.asm Resource (12680224)	26
Adding Main Adapter to IPMP Group Removes DNS Configuration (15816288)	26
Zone Does Not Boot if pkg:/system/resource-mgmt/resource-cap Is Not Installed and capped-memory Is Configured (15740089)	27
Active:Active ZFS Storage Appliance Clustered Configurations Are Not Supported (15521899)	27
Installation Notes	29
Installation Issues	29
Cannot Use scinstall in a CLI to Set Up Automated Installer Using Oracle Solaris Unified Archives (19058049)	29
Cannot Use autodiscovery to Discover the Transport Paths If Using VNICs That Are Created on the Same Physical Adapter (18245335)	29
scinstall Fails While Using Only IPv6 for a Public Network Configuration (16355496)	30
Hard to Determine Data Service Names for solaris10 Branded Zone Noninteractive Data Service Installation (15804349)	30
cacao Cannot Communicate on Machines Running Trusted Extensions (15803432)	31
Administration Notes	33
Administration Issues	33
-n Option Does Not Work With the clsnmp* Commands in Oracle Solaris Cluster 4.2 (19211571)	33
admin Property Does Not Work in clzonecluster (19223970)	33
Unable to Change the Zone Cluster's Configuration When Some of the Cluster Nodes are Down (19182915)	34
clresourcetype register Command Fails in the solaris10 Brand of Zone Clusters (19043805)	34
Unable To Reset acfs_mountpoint Property When the SUNW.scalable_acfs_proxy Resource Is Disabled (16928490)	35
clzonecluster install -a archive-no-cluster-pkgs zone-cluster Does Not Install Cluster Packages (18714803)	35
When Upgrading a Resource, Set the type_version Property First Before Updating Other Properties (18313716)	35
IPv6 Scalable Service Support Is Not Enabled By Default (15290321)	36
Removing a Node From an Exclusive-IP Zone Cluster Panics the Cluster Nodes (15817184)	36

Attempt to Delete Resource Group in Global Zone Fails Incorrectly (16704170)	37
LDom Stop Timeout in SUNWscxvm.stop Prevents LDom Unbound (18335346)	37
Allowing User to Remove Last DID Device From a Rawdisk Device Group and Later cldevicegroup Commands Result in cl_ccrad Core Dumps (15815165)	38
Public Net Failure Does Not Fail Over DB Server Resource With SCAN Listener (16231523)	38
EMC SRDF Rejects Switchover When Replicated Device-Group Status Will Cause Switchover and Switchback to Fail (15538295)	39
Cannot Create a Resource From a Configuration File With Non-Tunable Extension Properties (15658184)	39
Runtime Notes	41
Runtime Issues	41
clcomm: path online, Then Path Being Drained Every Eleven Seconds (18827672)	41
When a Shared-Address Resource Is Created Using the clressharedaddress Command or a Cluster Node Hosting the Shared-Address Resource Is Rebooted, a Syslog Error, "unable to get device major number for driver" Is Generated (15765209)	41
metaset -C release Hangs During shutdown -g0 (17038844)	42
clnode status -m -Z all Resulted in Core (18707228)	43
Graphical User Interface (GUI) Notes	45
GUI Issues	45
Autodiscovery Reports Error When Adding a Node to a Cluster (19173124)	45
GUI Does Not Refresh Data for Protection Groups in Certain Scenarios (19027066)	45
Refresh Issues in Pages Related to Zone Clusters (19144804)	45
Resource Creation Using the Resource Wizard in the GUI Sometimes Fails With a NullPointerException (19130707)	46
At Installation, cluster/manager:default Fails Fatally and Transitions to Maintenance Mode (18616858)	46
Property Editing: New Values Take One Refresh Cycle to Appear After Being Saved (18791809)	46
Unable to Install Oracle Solaris Cluster for solaris10 Branded Zone (19064831)	47

No Resource Types Listed in the Pull-Down Menu When Creating a New Resource Through the GUI Wizard (19013853)	47
Resource Group Not Listed in the GUI When It Is in the Protection Group But Has Been Removed From the Cluster (18720984)	47
Very Long Wait Times When Navigating to Partnership and Protection Group If a Cluster in the Partnership Is Down (18961504)	48
Data Services Notes	49
Data Services Issues	49
Oracle Database/WLS Resource Fails to Come Online Due to Locking Issue (15713853)	49
Disabled LDom Resource Is Keeping the LDom in Bound State (16323723)	49
Register Scripts Need to Define resource_offline_restart Dependency (15700470)	50
NFS Resource and HASP Resources Fail in Oracle Solaris Cluster 4.2 on Oracle Solaris 11.1 and Oracle Solaris 11.2 With zfs_recv (17365301)	50
Data Service Configuration Wizards Do Not Support Storage Resources and Resource Groups for Scalable HAStoragePlus (15820415)	51
Scalable Applications Are Not Isolated Between Zone Clusters (15611122)	51
Running clnas_add or clnas_remove Command on Multiple Nodes at the Same Time Could Cause Problem (15533979)	51
Developer Environment Notes	53
Developer Environment Issues	53
num_zoneclusters Property Cannot Be Set in Non-cluster Mode (1852819)	53
Geographic Edition Notes	55
Geographic Edition Issues	55
geosite_create Fails With JMX timeout, Dangling Heartbeat is Left Behind After geosite_delete (19232931)	55
Switchover Proceeds and Start Reports Success When the SRDF Device Group Is in a Partitioned State (17250376)	56
Oracle Data Guard Module Incorrectly Flagging SUNW.oracle_server Dependencies in the Single Instance (15818725)	56
Data Guard Module Incorrectly Updates rg_affinities and Protection Group Validation Failures When Multiple rac_server_proxy Resource Groups Exist on the Cluster (18940623)	56
Geographic Edition Runs show configuration Frequently and Results in the Database Getting Requests Seconds Apart (18779118)	57

Collision Problems Should Be Flagged at Protection Group Creation Time (15801862)	57
If Operation on Any Protection Group Fails During a Multigroup Operation, Then the Multigroup Operation Errors Out Without Doing Any Pending Operations on Other Protection Groups in the Multigroup (18270451)	57
If Resource Group Is in Pending_online_blocked State, geopg start Continuously Tries to Start the Resource Group (16636516)	58
After Changing Timeout of a Protection Group, an Attempt to Change the Timeout of a Multigroup Containing That Protection Group Fails in Some Situations (18436909)	58
Disallow site create or site add-member When Site of This Name Already Exists on Members (18415502)	59
Doing geosite update remote-cluster site on a Cluster Does Not Replicate the Site's Multigroups That Are Present on the Remote Cluster onto the Local Cluster (18368896)	59
Documentation Notes	61
System Administration Guide	61
HA for SAP NetWeaver Guide	61
Man Pages	63
cldevicegroup(1CL)	63
cluster(1CL)	64

Using This Documentation

- **Overview** – Lists product features, requirements, and restrictions for Oracle Solaris Cluster software and describes open defects and other known problems.
- **Audience** – Experienced system administrators with extensive knowledge of Oracle software and hardware.
- **Required knowledge** – Knowledge of the Oracle Solaris operating system and of Oracle Solaris Cluster software, and expertise with the volume manager software that is used with Oracle Solaris Cluster software.

This document is not to be used as a planning or presales guide.

Product Documentation Library

Late-breaking information and known issues for this product are included in the documentation library at <http://www.oracle.com/pls/topic/lookup?ctx=E39579>.

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.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program web site at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Feedback

Provide feedback about this documentation at <http://www.oracle.com/goto/docfeedback>.

Oracle® Solaris Cluster 4.2 Release Notes

This document provides the following information about the Oracle Solaris Cluster 4.2 software:

- “What's New in the Software” on page 13
- “Former Features Not Included in the Oracle Solaris Cluster 4.2 Software” on page 17
- “Product Localization Notes” on page 17
- “Configuring Accessibility Support for Oracle Solaris Cluster Manager” on page 18
- “Compatibility Issues” on page 23
- “Installation Issues” on page 29
- “Administration Issues” on page 33
- “Runtime Issues” on page 41
- “GUI Issues” on page 45
- “Data Services Issues” on page 49
- “Geographic Edition Issues” on page 55
- “Documentation Notes”

For the latest information about supported products and product versions for this release, see the [Oracle Solaris Cluster 4 Compatibility Guide \(http://www.oracle.com/technetwork/server-storage/solaris-cluster/overview/solariscluster4-compatibilityguide-1429037.pdf\)](http://www.oracle.com/technetwork/server-storage/solaris-cluster/overview/solariscluster4-compatibilityguide-1429037.pdf).

What's New in the Software

This section highlights information for existing customers about new features in Oracle Solaris Cluster 4.2 software.

The Oracle Solaris Cluster 4.2 SRU 1 provides the following new features:

- **New Oracle Solaris Cluster Manager wizards** - A new Tasks component with additional wizards is added to the Oracle Solaris Cluster Manager GUI. From the Tasks panel, you can run configuration wizards for the following components:
 - HA for Oracle Database data service

- Logical hostname resources
- Highly available storage resources (HAStoragePlus)

This functionality requires a minimum of Oracle Solaris Cluster 4.2 SRU 1.

For more information about how to set up and launch the GUI in a browser, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

- **Geographic Edition management of Oracle Database running remotely in an Oracle Data Guard configuration** - A configuration running the Geographic Edition framework can be used to manage Oracle Database in an Oracle Data Guard configuration that is running on a remote system that is not running Oracle Solaris Cluster software.

This functionality requires a minimum of Oracle Solaris Cluster 4.2 SRU 1.

For more information about how to set up and launch the GUI in a browser, see [Chapter 1, “Replicating Data With Oracle Data Guard Software,”](#) in [“Oracle Solaris Cluster Geographic Edition Data Replication Guide for Oracle Data Guard”](#).

The Oracle Solaris Cluster 4.2 software provides the following new features:

- **Disaster recovery orchestration** – Disaster recovery (DR) orchestration provides the ability to manage multiple Geographic Edition protection groups as a single entity known as a multigroup. This entity enables an administrator to coordinate disaster recovery protection of multiple clusters and multiple services on a per-site basis.

At initial release, this feature has the following configuration limitations:

- Maximum number of clusters in a site: 8
- Maximum number of protection groups in a multigroup: 10
- Maximum number of multigroups in a site: 10
- Maximum number of sites in a cluster: 8
- Maximum number of controllers in a site: 3

Contact your Oracle support representative to learn whether an increase of any of these limits has become qualified at a later date.

For more information, see [“Oracle Solaris Cluster Geographic Edition System Administration Guide”](#).

- **Oracle Solaris Cluster Manager Graphical User Interface (GUI)** – The Oracle Solaris Cluster Manager GUI helps you check the status of cluster and Geographic Edition components and perform administrative tasks for data services, nodes, resource groups, resources, NAS devices, adapters, cables, quorum devices, partnerships, protection groups, and so on. You can use the GUI in place of many of the command-line interface (CLI) commands. The GUI is installed when you install the Oracle Solaris Cluster 4.2 ha-cluster-full package.

For information about how to set up and launch the GUI in a browser, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

- **Secure Automated Installation for Oracle Solaris Cluster** – This feature provides secure automated installation for Oracle Solaris Cluster 4.2 by authenticating and encrypting/decrypting the communication between the Automated Installer (AI) server and the cluster nodes, as well as the installation of the IPS packages to the cluster nodes from the HTTPS repository locations.

For more information, see [“Establishing a New Oracle Solaris Cluster With the Automated Installer”](#) in [“Oracle Solaris Cluster Software Installation Guide”](#).

- **Oracle Solaris Unified Archive support in Oracle Solaris Cluster** – This feature enables you to use the AI server to install cluster nodes from Oracle Solaris Unified Archives. The new capabilities include:
 - Installing and configuring the cluster with the new configuration
 - Restoring cluster nodes from the recovery archives that are created for the specific nodes
 - Replicating a new cluster from an existing cluster that has a similar hardware configuration

On cluster nodes that are already running, you can use the `clzonecluster install` command to install a new zone cluster and the `clzonecluster configure` command to configure a new zone cluster from Oracle Solaris Unified Archives.

For more information, see [“Establishing a New Oracle Solaris Cluster With the Automated Installer”](#) in [“Oracle Solaris Cluster Software Installation Guide”](#).

- **Enhancements to the Oracle Solaris Cluster APIs to provide improved support for HA data services that perform live migration** – The Resource Group Manager (RGM) supports a new resource property, `Pre_evict`, and a new `scha_resourcegroup_get` query tag, `SCHA_TARGET_NODES`, for the use of data services developers. These new API features allow a data service that performs a live migration kind of switchover from the `Stop` method to find out the target node of a switchover, and to trigger any required resource group evictions before the switchover is initiated.

An example of one such data service is HA for Oracle VM for SPARC. The `TARGET_NODES` query now allows live migration to be used in several other instances, such as `scha_control` giveovers, in which live migration was previously not possible. Pre-eviction enables Oracle Solaris Cluster to move excess workload off the target node before the switchover begins, which causes the live migration to succeed more often. The HA for Oracle VM for SPARC data service is enhanced to take advantage of these new API features.

For more information, see [“Supporting Resource Types That Perform Resource Migration From Their Stop Method”](#) in [“Oracle Solaris Cluster Data Services Developer’s Guide”](#).

- **Enhancement to the Cluster Event SNMP Interface** – This feature enhances the `clsnmpmib` utility by introducing a new severity, `NOTICE`. This new severity and other higher severity events now have a one-to-one correlation for the configuration and status change of cluster objects. With this feature, you can configure the minimum severity at which events will be logged in Management Information Base (MIB). You can also indicate the number of events to be logged in MIB.

For more information, see [“Creating, Setting Up, and Managing the Oracle Solaris Cluster SNMP Event MIB”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

- **Support for Oracle Solaris Kernel Zones** – In this release, the Oracle Solaris Cluster Data Service for Oracle Solaris Zones supports Oracle Solaris Kernel Zones (solaris-kz brand). The sczbt component now supports cold and warm migration for the kernel zones. The three data service components are implemented as their own dedicated resource types as follows:
 - sczbt – ORCL.ha-zone_sczbt
 - sczsh – ORCL.ha-zone_sczsh
 - sczsmf – ORCL.ha-zone_sczsmf

Resource configuration is still done by amending the component configuration file and providing it to the component register script.

For more information, see [“Oracle Solaris Cluster Data Service for Oracle Solaris Zones Guide”](#).

- **Support for a new version of Generic Data Service (GDSv2)** – GDSv2 includes new generic resource types and more functionality than the previous version of GDS.
For more information, see [Chapter 2, “Creating a Data Service with GDSv2,”](#) in [“Oracle Solaris Cluster Generic Data Service \(GDS\) Guide”](#).
- **Support for Oracle Solaris Cluster HA for Oracle JD Edwards EnterpriseOne Enterprise Server (SPARC Only)** – This high availability data service for the JD Edwards EnterpriseOne application enables the Oracle Solaris Cluster product to start, stop, and monitor the Enterprise Server component of JD Edwards EnterpriseOne product. The new resource type interrogates the JD Edwards EnterpriseOne Enterprise Server processes and interprets the availability of those processes as an Oracle Solaris Cluster resource state or status in an Oracle Solaris Cluster configuration.

For more information, see [“Oracle Solaris Cluster Data Service for Oracle JD Edwards EnterpriseOne Enterprise Server Guide”](#).

- **Support for Oracle Solaris Cluster HA for Oracle GoldenGate** – In this release, this high availability data service for Oracle GoldenGate supports the Oracle GoldenGate versions 11.2.1.x and 12.1.2.x.

For more information, see [“Oracle Solaris Cluster Data Service for Oracle GoldenGate Guide”](#).

- **Support for Oracle 12c RAC Database Containers** – The Oracle Real Application Clusters (RAC) data service now supports the database containers feature of Oracle RAC 12c.

For more information, see [“Oracle Solaris Cluster Data Service for Oracle External Proxy Guide”](#).

- **Support for Oracle RAC Database Services Agent** – This feature enables Oracle RAC database services to be represented by proxy resources in Oracle Solaris Cluster, thereby allowing application dependencies to be fine-tuned for better availability.

For more information, see [“Oracle Solaris Cluster Data Service for Oracle External Proxy Guide”](#).

- **Support for policy-managed RAC databases** – This feature allows the use of policy-managed RAC databases with the scalable RAC server proxy resource type.

For information about this database feature, see your Oracle Database documentation. For more information on configuring Oracle Solaris Cluster for this feature, see [“Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide”](#).

Former Features Not Included in the Oracle Solaris Cluster 4.2 Software

The following features are included in the Oracle Solaris Cluster 3.3 version but are not included in the Oracle Solaris Cluster 4.2 release:

- Support for Veritas File System (VxFS) and Veritas Volume Manager (VxVM)
- Support for the VxVM cluster feature for Oracle RAC in addition to VxVM with Oracle Solaris Cluster
- Support for non-global zones as resource-group node-list targets
- Support for Oracle Solaris IP Security Architecture (IPsec) on the private interconnect
- Support for Hitachi True Copy and Hitachi Universal Replicator storage-based replication

Product Localization Notes

Localization for certain components of Oracle Solaris Cluster 4.2 software is as follows:

- Software command line - Japanese, Korean, Simplified Chinese
- Software GUI - French, Japanese, Korean, Simplified Chinese, Spanish
- Online help - French, Japanese, Korean, Simplified Chinese, Spanish
- Man pages - Japanese, Simplified Chinese

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

Shell	Command
sh	\$ LC_MESSAGES=C;export LC_MESSAGES
ksh	\$ export LC_MESSAGES=C
bash	\$ export LC_MESSAGES=C
csh	% setenv LC_MESSAGES C

Shell	Command
tcsh	% setenv LC_MESSAGES C

Configuring Accessibility Support for Oracle Solaris Cluster Manager

The Oracle Solaris Cluster Manager graphical user interface (GUI) is built using Oracle Application Development Framework (ADF) Faces. See <http://www.oracle.com/us/corporate/accessibility/templates/t2-3281.html> for more details about ADF Faces.

Two elements can be configured in the `trinidad-config.xml` file to enable accessibility features: `<accessibility-mode>` and `<accessibility-profile>`. By default, neither element is configured. This section describes the available accessibility options and how to configure them.

`<accessibility-mode>` Element

Oracle Solaris Cluster Manager provides two levels of application-level accessibility support configured in the `trinidad-config.xml` file using the `<accessibility-mode>` element. The acceptable values for `<accessibility-mode>` are as follows:

default

By default, ADF Faces generates components that have rich user interface interaction and are also accessible through the keyboard. Note that, in the `default` mode, screen readers cannot access all ADF Faces components.

```
<accessibility-mode>default</accessibility-mode>
```

screenReader

ADF Faces generates components that are optimized for use with screen readers. The `screenReader` mode facilitates the display for visually impaired users but will degrade the display for users without visual impairment.

```
<accessibility-mode>screenReader</accessibility-mode>
```

`<accessibility-profile>` Element

You can specify the `<accessibility-profile>` element in the `trinidad-config.xml` file. Valid options are `high-contrast`, `large-fonts`, or `both`. The acceptable values for `<accessibility-profile>` are as follows:

high-contrast

ADF Faces can generate high-contrast–friendly visual content. The ADF Faces high-contrast mode makes ADF Faces applications compatible with operating systems or browsers that have high-contrast features enabled. For example, ADF Faces changes its use of background images and background colors in high-contrast mode to prevent the loss of visual information.

The ADF Faces high-contrast mode is more beneficial if used in conjunction with your browser's or operating system's high-contrast mode. Also, some users might find it beneficial to use `large-fonts` mode along with high-contrast mode.

```
<accessibility-profile>high-contrast</accessibility-profile>
```

large-fonts

ADF Faces can generate browser zoom–friendly content. In default mode, most text and many containers have a fixed font size to provide a consistent and defined look. In `large-fonts` mode, text and containers have a scalable font size, which enables ADF Faces to be compatible with browsers that are set to larger font sizes and to work with browser zoom capabilities.

Some users might find it beneficial to use high-contrast mode along with `large-fonts` mode. However, if you are not using a large-font mode or browser-zoom capabilities, disable `large-fonts` mode.

```
<accessibility-profile>large-fonts</accessibility-profile>
```

If you require both the high-contrast and `large-fonts` options, include the following entry in the `trinidad-config.xml` file:

```
<accessibility-profile>high-contrast large-fonts</accessibility-profile>
```

If application-level access requires all the options to be set, include the following entry in the `trinidad-config.xml` file:

```
<accessibility-mode>screenReader</accessibility-mode>
<accessibility-profile>high-contrast large-fonts</accessibility-profile>
```

▼ How to Configure Accessibility Support in the `trinidad-config.xml` File

This procedure describes how to manually edit the `trinidad-config.xml` file to set accessibility modes for the Oracle Solaris Cluster Manager GUI. This file is located in the following directory path:

```
/var/cluster/lib/ClusterManager/glassfish3/glassfish/domains/domain1/applications\
/ClusterManager_GF/Yosemite_View_Glassfish_webapp_war/WEB-INF/trinidad-config.xml
```

Tip - The configuration changes made in this procedure operate at a global level on a given cluster node. This procedure alters application behavior for all users of the GUI connecting through a browser as opposed to those authenticating into that particular cluster node. To accommodate both GUI users who need assistive technology and those who do not, apply these settings to some but not all of the cluster nodes.

For example, in a two-node cluster, you would configure the pnode-1 node with the accessibility options but not configure them on the pnode-2 node. Users who do not need assistive technology would then point their browsers at pnode-2, while those who do need it would use pnode-1.

In either case, a user can authenticate into whichever node they want to access by specifying it in the Cluster Node field on the login page.

1. **Assume the root role.**
2. **Disable the Oracle GlassFish server.**

```
# svcadm disable -s manager-glassfish3
```
3. **Open the trinidad-config.xml file for editing.**
4. **Add entries to the trinidad-config.xml file to set the desired accessibility options.**
See [“Configuring Accessibility Support for Oracle Solaris Cluster Manager” on page 18](#) for descriptions of accessibility options and the trinidad-config.xml entry to use for the different options or combinations of options.
5. **Enable the Oracle GlassFish server.**

```
# svcadm enable -s manager-glassfish3
```

Example 1 trinidad-config.xml File With Accessibility Configured

The following example shows a trinidad-config.xml file configured with all accessibility options.

```
<?xml version="1.0" encoding="UTF-8"?>
<trinidad-config xmlns="http://myfaces.apache.org/trinidad/config">
<skin-family>skyros</skin-family>
<skin-version>v1</skin-version>
<!--
<skin-family>fusionFx</skin-family>
<skin-version>v2</skin-version>
-->
<accessibility-mode>screenReader</accessibility-mode>
```

```
<accessibility-profile>high-contrast large-fonts</accessibility-profile>  
</trinidad-config>
```

Commands Modified in This Release

The following commands have been modified in this release.

- `cldevicegroup create` - You cannot create a device group without specifying at least one device.
- `cldevicegroup remove-device` - You cannot use the `remove-device` subcommand to remove all the devices in a device group. A device group must contain at least one device. To remove all the devices in a device group, instead use the `delete` subcommand to remove the device group.

Compatibility Notes

This chapter contains information about Oracle Solaris Cluster compatibility issues with other products. Contact Oracle support services to find out whether a code fix has become available.

Compatibility Issues

Restriction of Support for Solaris Volume Manager Multi-Owner Disk Sets

Oracle Solaris Cluster 4.2 software does not support Solaris Volume Manager multi-owner disk sets. This restriction of support is due to multiple Oracle Solaris bugs that might adversely affect cluster availability.

Until support of Solaris Volume Manager multi-owner disk sets in Oracle Solaris Cluster 4.2 configurations can be qualified, documentation for use of multi-owner disk sets should be ignored.

Contact your Oracle support representative to learn whether Solaris Volume Manager multi-owner disk sets becomes supported for Oracle Solaris Cluster 4.2 configurations.

Creation of labeled Brand of Zone Clusters Fail (19074503)

Problem Summary: Creation of the labeled brand of zone clusters fails on Oracle Solaris 11.2.

Workaround: Add the following lines to the `/usr/lib/brand/labeled/config.xml` file.

```
<resource name="zone">
<property name="pool" reset-on-clone="false"/>
</resource>
```

The three lines added to the `/usr/lib/brand/labeled/config.xml` file must match the same three lines in the `/usr/lib/brand/solaris/config.xml` file.

Following is an example of the `config.xml` file for the labeled brand of zone cluster.

```
<brand name="labeled" supports_zonestatd="true" requires_hostid="false">
...
...
<resource name="zone">
<property name="pool" reset-on-clone="false"/>
</resource>
</brand>
```

show configuration Attempts Pile Up and Lead to ORA-00020 (18746643)

Problem Summary: A `show configuration` command on the standby cluster hangs especially when there are network connection issues between the primary cluster and the standby cluster. As a result, `show configuration` requests pile up sometimes, leading to maximum number of processes exceeded (ORA-00020).

Workaround: Apply the Oracle Database patch 18746643.

Cacao Creates a Huge File and Host Runs Out of Root Space Quickly (18099260)

Problem Summary: This problem occurs when cacao fails several times, triggering a watchdog recovery action that launches a script that runs `jstack`. If the version of `jstack` does not match the Java version, the resulting mismatch errors are sent to the `/var/cacao/instances/default/proc_scan.txt` file until the file system is full.

Workaround: Ensure that the cluster is running at least Oracle Solaris 11.1 SRU 19.

The Node Panics With FreeNode Message While Starting the Oracle Solaris Clusterware Inside the solaris10 Brand of Zone Cluster (19032049)

Problem Summary: You might see this issue in the Oracle Real Application Cluster (RAC) releases configured to use the Oracle Solaris Cluster Native SKGXN functionality in a `solaris10` brand of zone cluster. The system panics either while installing or starting the Cluster Ready Services (CRS). This problem occurs when the global cluster nodes are running on Oracle Solaris 11.2.

Workaround: Contact Oracle support services to see if a patch or a workaround is available.

VNIC Names Longer Than 16 Characters Cause Problems (17362337)

Problem Summary: If you use long names for VNICs in exclusive-IP zone clusters (solaris and solaris10 brands of zone clusters), you might not be able to choose the VNIC during system configuration.

Workaround: When using VNICs for zone clusters, the name of the VNIC must be less than 16 characters long.

sdplib: ibt_get_ip_paths With APM Failed Status 25 on Oracle Solaris Cluster When Port Is Down (18692783)

Problem Summary: In Oracle Solaris 11.2, the Sockets Direct Protocol (SDP) that uses the Oracle Solaris Cluster's private network might fail when one redundant physical IB network path or port is faulted.

By default, Oracle Solaris 11.2 supports HCA Managed Failover based on IPMP and APM for SDP traffic. In an Oracle Solaris Cluster environment, the interconnects are not managed under IPMP and a transparent automatic recovery of the SDP traffic upon InfiniBand link failure is required.

Workaround: To enable the automatic recovery capability and to disable the HCA Managed Failover support, add the following line in the /etc/system file:

```
set sdplib:sdp_ip_path_verify=0
```

Reboot the node with the boot archive updated for the change to be effective.

Zone Cluster Does Not Recognize *dsk WildCard With svm /dev/md (18167722)

Problem Summary: If you attempt to add Solaris Volume Manager metadevices to a zone cluster node using the asterisk "*" as a wildcard, the device will not be added. For example:

```
clzc:zc_name> add device
clzc:zc_name:device> set match=/dev/md/shared/*dsk/d11
clzc:zc_name:device> end
```

Workaround: Use explicit path names as shown in the following example.

```
clzc:zc_name:device> set match=/dev/md/shared/dsk/d11
```

```
clzc:zc_name:device> end
clzc:zc_name:device> set match=/dev/md/shared/rdisk/d11
clzc:zc_name:device> end
```

If a device with the wildcard designation has been erroneously added, it can be removed as shown in the following example.

```
clzc:zone_name> remove device match=/dev/md/shared/2/*disk/d11
```

osysmond Core Dumps in Solaris10 Branded Zone During GI root.sh and Starting of CRS (14456069)

Problem Summary: This problem involves Oracle RAC 11g release 2 configured in a solaris10 brand of zone cluster. When the Grid Infrastructure root.sh script is run or when Cluster Ready Services (CRS) is started, the osysmond process might dump core one or more times.

Workaround: Contact Oracle Support to learn whether a patch or workaround is available.

Oracle Clusterware Fails to Create All SIDs for ora.asm Resource (12680224)

Problem Summary: When creating an Oracle Solaris Cluster resource for an Oracle ASM instance, one of the following error messages might be reported by the clsetup utility:

```
ORACLE_SID (+ASM2) does not match the Oracle ASM configuration ORACLE_SID () within CRS
ERROR: Oracle ASM is either not installed or the installation is invalid!
```

This situation occurs because, after Oracle Grid Infrastructure 11g release 2 is installed, the value for GEN_USR_ORA_INST_NAME@SERVERNAME of the ora.asm resource does not contain all the Oracle ASM SIDs that are running on the cluster.

Workaround: Use the crsctl command to add the missing SIDs to the ora.asm resource.

```
# crsctl modify res ora.asm \
-attr "GEN_USR_ORA_INST_NAME@SERVERNAME(hostname)"=ASM_SID
```

Adding Main Adapter to IPMP Group Removes DNS Configuration (15816288)

Problem Summary: A problem occurs if you delete a network adapter and then re-create it for an IPMP group, such as in the following example commands:

```
# ipadm delete-ip adapter
# ipadm create-ip adapter
# ipadm create-ipmp -i adapter sc_ipmp0
# ipadm create-addr -T static -a local=hostname/24 sc_ipmp0/v4
```

Soon after the IPMP address is created, the `/etc/resolv.conf` file disappears and the LDAP service becomes disabled. Even an enabled service stays in the offline state.

Workaround: Before you delete the network adapter with the `ipadm delete-ip` command, run the `svcadm refresh network/location:default` command.

Zone Does Not Boot if pkg:/system/resource-mgmt/resource-cap Is Not Installed and capped-memory Is Configured (15740089)

Problem Summary: If the package `pkg:/system/resource-mgmt/resource-cap` is not installed and a zone is configured with the `capped-memory` resource control as part of the configuration, the zone boot fails. Output is similar to the following:

```
zone 'zone-1': enabling system/rcap service failed: entity not found
zoneadm: zone 'zone-1': call to zoneadmd failed
```

Workaround: Install the `pkg:/system/resource-mgmt/resource-cap` package into the global zone. Once the `resource-cap` package is installed, the zone can boot.

Active:Active ZFS Storage Appliance Clustered Configurations Are Not Supported (15521899)

Problem Summary: Simultaneously replicating from both heads in an active:active clustered ZFS SA configuration is not supported. This is a ZFS Storage Appliance product restriction (see Bug 15521899).

Workaround: Active:passive configurations are currently supported in a clustered configuration.

Installation Notes

This chapter contains known issues and bugs that affect the installation of Oracle Solaris Cluster 4.2. Contact Oracle support services to find out whether a code fix has become available.

Installation Issues

Cannot Use `scinstall` in a CLI to Set Up Automated Installer Using Oracle Solaris Unified Archives (19058049)

Problem Summary: Using the `scinstall` command to set up Automated Installer to install cluster nodes from Oracle Solaris Unified Archives does not set it up properly.

Workaround: Use the interactive `scinstall` by running `/usr/cluster/bin/scinstall` without any options. Then, choose the options from the menus and provide your input.

Cannot Use autodiscovery to Discover the Transport Paths If Using VNICs That Are Created on the Same Physical Adapter (18245335)

Problem Summary: When configuring a cluster, if you use VNICs that are created on the same physical adapter as the transport adapters, using autodiscovery while running `scinstall` to find the transport paths will fail.

Workaround 1: Instead of using interactive `scinstall` to configure a cluster, configure the nodes to join the cluster one by one and do not use autodiscovery.

Workaround 2: Do not use the interactive `scinstall` method. Instead, use the CLI and specify the virtual adapters accordingly.

scinstall Fails While Using Only IPv6 for a Public Network Configuration (16355496)

Problem Summary: `scinstall` fails while configuring a cluster if the public network has only IPv6 addresses configured. You might get an error message stating that the second node is an unknown host.

Workaround: Configure IPv4 addresses just for the `scinstall` configuration. Once the cluster has been formed, the IPv4 addresses can be removed.

Hard to Determine Data Service Names for solaris10 Branded Zone Noninteractive Data Service Installation (15804349)

Problem Summary: Determining the agent names to specify when using the `clzonecluster install-cluster` command to install agents with the `-s` option is difficult.

Workaround: When using the `clzonecluster install-cluster -d dvd -s {all | software-component[,...]} options zone-cluster` command to create a `solaris10` brand of zone cluster, you can specify the following cluster components with the `-s` option:

- `geo`
- `9ias`
- `apache`
- `container`
- `dhcp`
- `dns`
- `ebs (SPARC only)`
- `hadb`
- `ids`
- `iws`
- `kerberos`
- `livecache`
- `mqi`
- `mqs`
- `mys`
- `n1ge`
- `n1sps`

- nfs
- obiee (SPARC only)
- oep
- ohs
- opmn
- oracle
- pax (SPARC only)
- PeopleSoft (SPARC only)
- PostgreSQL
- rac
- slas
- slmq
- saa (SPARC only)
- sag (SPARC only)
- sap
- sapdb
- sapnetw
- sapwebas
- siebel (SPARC only)
- smb
- sybase
- TimesTen
- tomcat
- wls
- xvm (SPARC only)

cacao Cannot Communicate on Machines Running Trusted Extensions (15803432)

Problem Summary: If the Trusted Extensions feature of Oracle Solaris software is enabled before the Oracle Solaris Cluster software is installed and configured, the Oracle Solaris Cluster setup procedures are unable to copy the common agent container security keys from one node to other nodes of the cluster. Identical copies of the security keys on all cluster nodes is a requirement for the container to function properly on cluster nodes.

Workaround: Manually copy the security keys from one global-cluster node to all other nodes of the global cluster.

1. On each node, stop the security file agent.

```
phys-schost# /usr/sbin/cacaoadm stop
```

2. On one node, change to the `/etc/cacao/instances/default/` directory.

```
phys-schost-1# cd /etc/cacao/instances/default/
```

3. Create a tar file of the `/etc/cacao/instances/default/` directory.

```
phys-schost-1# tar cf /tmp/SECURITY.tar security
```

4. Copy the `/tmp/SECURITY.tar` file to each of the other cluster nodes.

5. On each node to which you copied the `/tmp/SECURITY.tar` file, extract the security files.

Any security files that already exist in the `/etc/cacao/instances/default/` directory are overwritten.

```
phys-schost-2# cd /etc/cacao/instances/default/
```

```
phys-schost-2# tar xf /tmp/SECURITY.tar
```

6. Delete the `/tmp/SECURITY.tar` file from each node in the cluster.

Note - You must delete each copy of the tar file to avoid security risks.

```
phys-schost-1# rm /tmp/SECURITY.tar
```

```
phys-schost-2# rm /tmp/SECURITY.tar
```

7. On each node, restart the security file agent.

```
phys-schost# /usr/sbin/cacaoadm start
```


Administration Notes

This chapter contains known issues and bugs that affect the administration of Oracle Solaris Cluster 4.2. Contact Oracle support services to find out whether a code fix has become available.

Administration Issues

-n Option Does Not Work With the `clsnmp*` Commands in Oracle Solaris Cluster 4.2 (19211571)

Problem Summary: Attempting to use `-n nodelist` in the `clsnmp*` commands will fail.

Workaround: Do not use the `-n` option to specify the nodes. Instead, run the desired command on each node, one at a time.

`admin` Property Does Not Work in `clzonecluster` (19223970)

Problem Summary: If the `admin` resource is present in the zone cluster configuration, the `clzonecluster info` command shows a null value or a garbage value for the `auths` property of the `admin` resource. If you attempt to change or add a new `admin` resource to the zone cluster configuration, the zone cluster creation or update will fail. Updating the zone cluster might incorrectly change the zone configuration and result in the `clzonecluster status` command reporting the zone status to be in the unknown state.

Workaround: If an `admin` resource has to be added or an existing `admin` resource has to be modified, use the `solaris zonecfg` command to add or modify the `admin` property for each zone of the zone cluster. Do not use the `clzonecluster` command to add or modify an `admin` resource.

If you changed a zone cluster configuration using the `clzonecluster` command, then the admin resource in the zone configuration will not exist and has to be added back to the zone configuration.

Unable to Change the Zone Cluster's Configuration When Some of the Cluster Nodes are Down (19182915)

Problem Summary: This problem might occur when you attempt to update the zone cluster's configuration when some of the cluster nodes hosting the zone cluster are down. It occurs in the `solaris10` and `labeled` brands of zone clusters. It also occurs in the `solaris` brand of zone clusters when the admin resource is configured.

Workaround: Update the zone cluster when all the nodes hosting the zone cluster are up and part of the cluster membership. If it is not possible to get the nodes up, update the configuration using the `solaris zonecfg` command.

However, when the nodes are up, perform the same update using the `clzonecluster` command, so that the subsequent update using the `clzonecluster` command does not wipe out the update performed using the `solaris zonecfg` command.

`clresourcetype register` Command Fails in the `solaris10` Brand of Zone Clusters (19043805)

Problem Summary: Sometimes the `clresourcetype register` command might fail when run on a `solaris10` brand of zone clusters. You might see one of the following errors:

```
# clrt register -f /SUNW.test_rt_deps -p RT_SYSTEM=TRUE RT1_spd_6
clrt: (C549566) Failed to find "zc2:RT1_spd_" in RTR file
"/zones/zc2/root/SUNW.test_rt_deps"

# clrt register -N SUNW.gds
clrt: (C637184) Unrecognized option - "- ".
clrt: (C637184) Unrecognized option - "-S".
clrt: (C637184) Unrecognized option - "-U".
clrt: (C637184) Unrecognized option - "-W".
clrt: (C637184) Unrecognized option - "-.".
clrt: (C637184) Unrecognized option - "-g".
clrt: (C637184) Unrecognized option - "-d".
clrt: (C637184) Unrecognized option - "-s".
clrt: (C425781) Usage errors.
```

```
Usage: clrt register [<options>] + | <resourcetype> ...
```

```
# geopg get -s partnershipname
```

Operation failed for following protection groups:
Unable to register resource type (0).

Workaround: Register the resource type in the global zone.

```
# clrt register -Z zone name resource type
```

Note - In the Oracle Solaris Cluster Geographic Edition, it will not be possible to use the workaround when a command such as `geopg get` fails. Contact Oracle support services to see if a patch is available.

Unable To Reset `acfs_mountpoint` Property When the `SUNW.scalable_acfs_proxy` Resource Is Disabled (16928490)

Problem Summary: You cannot reset the `acfs_mountpoint` property of a `SUNW.scalable_acfs_proxy` resource if it is disabled.

Workaround: The `acfs_mountpoint` property can be set only at resource creation time. To change the `acfs_mountpoint` property, you must delete and re-create the resource.

`clzonecluster install -a archive-no-cluster-pkgs zone-cluster` Does Not Install Cluster Packages (18714803)

Problem Summary: This issue might occur when you install a zone cluster from an Oracle Solaris Unified Archive created from a source that does not have cluster packages in it.

Workaround: Manually install the cluster packages within the zone cluster.

When Upgrading a Resource, Set the `type_version` Property First Before Updating Other Properties (18313716)

Problem Summary: When you upgrade a resource to a newer version, setting the `type_version` property and other properties in the same command will fail.

Workaround: Instead of changing `type_version` and other properties in the same `clresource set` command, change the `type_version` property first and then change other properties in a separate `clresource set` command.

IPv6 Scalable Service Support Is Not Enabled By Default (15290321)

Problem Summary: IPv6 interfaces are not plumbed on the interconnect adapters by default. This IPv6 plumbing is required for forwarding IPv6 scalable service packets.

Workaround: All cluster nodes must first be prepared to run IPv6. This preparation includes proper configuration of network interfaces, server/client application software, name services, and routing infrastructure. Not doing so could result in unexpected failures of network applications. See your Oracle Solaris documentation on IPv6 before enabling IPv6 scalable services on a cluster.

To enable IPv6 scalable service support:

1. Add the following line to `/etc/system` on all nodes.

```
set cl_comm:ifk_disable_v6=0
```

2. Enable IPv6 plumbing.

- If reboot is allowed, reboot all cluster nodes. Reboot the nodes one at a time to minimize outage time.
- If reboot is not allowed, run the following utility to enable IPv6 plumbing on the interconnect adapters.

```
# /usr/cluster/lib/sc/config_ipv6
```

This utility brings up an IPv6 interface on all the cluster interconnect adapters with a link-local address. It enables proper forwarding of IPv6 scalable service packets over the interconnects.

Removing a Node From an Exclusive-IP Zone Cluster Panics the Cluster Nodes (15817184)

Problem Summary: When a zone-cluster node is removed from an exclusive-IP zone cluster, the global-cluster nodes that host the exclusive-IP zone cluster panics. The issue is seen only on a global-cluster with InfiniBand interconnects.

Workaround: Halt the exclusive-IP zone cluster before you remove the zone-cluster node.

Attempt to Delete Resource Group in Global Zone Fails Incorrectly (16704170)

Problem Summary: When both the global cluster and a zone cluster contain a resource group by the same name, if either resource group has a non-empty value for the `RG_dependencies` or `RG_affinities` property, then attempting to remove the other resource group with the same name will fail with a complaint about the existing dependency.

Workaround: Include the `-F` (force) option in the command to delete the resource group. For example:

```
# clrg delete -F rg-name
```

Then, recreate any dependencies or affinities on the remaining resource group with the same name if deleted with the force option.

LDom Stop Timeout in SUNWscxvm.stop Prevents LDom Unbound (18335346)

Problem Summary: If `ldm stop` is timed out in `SUNWscxvm.stop`, the logical domain (LDom) remains bounded and prevents the cluster resource group from cleanly shutting down the domain. This behavior prevents the successful failover of the resource group. Also, currently the `STOP TIMEOUT` value is not taken into consideration while stopping the LDom and it automatically time outs after 60 seconds.

You might encounter one of the following error messages.

```
[ID 885590 daemon.notice] Domain domain_name has been forcefully terminated.
[ID 567783 daemon.notice] domain stop result code : 0 - ldom_name stop timed out. The domain
might still be in the process of shutting down.
[ID 567783 daemon.notice] domain stop result code : 0 - Either let it continue, or specify -f
to force it to stop.
[ID 567783 daemon.notice] domain stop result code : 0 - LDom ldom_name cannot be unbound
because it is stopping
[ID 567783 daemon.notice] domain stop result code : 0 - LDom ldom_name stopped
```

When you run `ldm list`, you will see that the LDom is in a bound state.

```
# ldm list
NAME STATE FLAGS CONS VCPU MEMORY UTIL NORM UPTIME
primary active -n-cv- UART 16 15872M 0.1% 0.1% 28m
ldom-1 active -n---- 5000 8 8G 0.0% 0.0% 28m
ldom-2 bound ----- 5001 112 112G
```

When the `clresourcegroup online` command is issued after the LDom stop timeout messages, it hangs because it has been forcibly terminated.

Workaround: This workaround works only for nodes with at least Logical Domains Manager 3.1. To modify the LDom timeout value:

Edit the `/opt/SUNWscxvm/bin/functions` file on all the nodes by replacing:

```
 ${HATIMERUN} -t ${MAX_STOP_TIMEOUT} -k KILL ${LDM} stop-domain ${DOMAIN} >> $LOGFILE 2&1
```

with:

```
 LDOM_TIMEOUT=$((MAX_STOP_TIMEOUT*80/100)) ${HATIMERUN} -t ${MAX_STOP_TIMEOUT} -k KILL ${LDM}
 stop-domain -t ${LDOM_TIMEOUT} ${DOMAIN} >> $LOGFILE 2&1
```

Now the `ldm stop` timeout is not 60 seconds but `LDOM_TIMEOUT` seconds. The `ldm stop-domain` command is issued first if the LDom does not shutdown within `LDOM_TIMEOUT` seconds. Then `ldm stop-domain -q` is issued. This command automatically forces the LDom to stop. To allow time for `ldm stop-domain -q` to execute, `LDOM_TIMEOUT` has been set to 80% of `MAX_STOP_TIMEOUT`.

Allowing User to Remove Last DID Device From a Rawdisk Device Group and Later `cldevicegroup` Commands Result in `cl_ccrad` Core Dumps (15815165)

Problem Summary: In a cluster that has multiple device groups, if a busy device is removed from one of the device groups leaving that group empty, all the other device groups might not be able to go offline. Subsequent `cldevicegroup` commands might cause a core dump.

Workaround: Do not remove a busy device from a device group. Do not remove the last device from a device group.

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.

EMC SRDF Rejects Switchover When Replicated Device-Group Status Will Cause Switchover and Switchback to Fail (15538295)

Problem Summary: If 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 has been returned to a paired state.

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

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.

Runtime Notes

This chapter contains known issues and bugs that affect the runtime of Oracle Solaris Cluster 4.2. Contact Oracle support services to find out whether a code fix has become available.

Runtime Issues

clcomm: path online, Then Path Being Drained Every Eleven Seconds (18827672)

Problem Summary: On an InfiniBand cluster, when a faulty network link comes back online, occasionally the path gets drained and re-created.

Workaround: Use the `clintr` command to disable the cable path and then re-enable the path.

When a Shared-Address Resource Is Created Using the clressharedaddress Command or a Cluster Node Hosting the Shared-Address Resource Is Rebooted, a Syslog Error, "unable to get device major number for driver" Is Generated (15765209)

Problem Summary: When a `SUNW.SharedAddress` resource is created using the `clrssa` command or a cluster node hosting the Shared-Address resource is rebooted, a syslog error, "unable to get device major number for driver" is generated. There is no impact on scalable services functionality as long as all the interfaces of the IPMP groups hosting the scalable address are plumbed. If an adapter that is part of the IPMP group is not plumbed, the autopush functionality to push the scalable services module on top of the adapter is disabled.

Workaround: Verify whether the adapter that is part of the IPMP group hosting the scalable address has the module `mcnet` in the stream of the device by executing the following command:

```
# /usr/sbin/ifconfig interface modlist
```

If not, manually push the `mcnet` module on top of the device by executing the following command so that `mcnet` is pushed on top of the device under the IP module.

```
# /usr/sbin/ifconfig interface modinsert mcnet@pos
```

metaset -C release Hangs During shutdown -g0 (17038844)

Problem Summary: Transitioning to the single-user runtime level from the multiuser level is not supported.

If this command is issued under certain configurations, shutdown of the node might hang, preventing evacuation (switchover) of services on that node.

Workaround: On the node where the shutdown is hanging, use either `halt(1M)` or `reboot(1M)` to force the node to shut down.

To prevent the hang from occurring, avoid executing any of the following commands for shutting down a cluster node:

- `init S`
- `init s`
- `init 1`
- `init 2`
- `shutdown [without -i flag]`

Instead, use one of the following commands:

- To take a node down to firmware:
`shutdown -i 0 or init 0`
- To take a node down and power off:
`shutdown -i 5 or init 5`
- To reboot a node:
`shutdown -i 6 or init 6`
- To shut down all cluster nodes to firmware, execute this command on one node:
`cluster shutdown`

clnode status -m -Z all Resulted in Core (18707228)

Problem Summary: Issuing the `clnode status -m -Z all exclusive-IP-zone-cluster-name` command on multiple nodes simultaneously, that is, from the master console, might cause a core dump of `clnode` and one or more nodes will hang.

Workaround: Do not run this command on multiple nodes simultaneously. If any node hangs, log in to the exclusive-IP zones on each node, kill the `pnm_mod_serverd` process, and restart it by running `/usr/cluster/lib/sc/pnm_mod_serverd`.

Graphical User Interface (GUI) Notes

This chapter contains known issues and bugs that affect the graphical user interface GUI of Oracle Solaris Cluster 4.2. Contact Oracle support services to find out whether a code fix has become available.

GUI Issues

Autodiscovery Reports Error When Adding a Node to a Cluster (19173124)

Problem Summary: While adding a new node to an existing cluster, autodiscovery of the appropriate transport adapter does not work.

Workaround: Do not choose autodiscovery to add a new node to an existing cluster. Instead, choose the appropriate transport adapters from the list of adapters offered.

GUI Does Not Refresh Data for Protection Groups in Certain Scenarios (19027066)

Problem Summary: Sometimes the GUI does not refresh the protection group data. As a result, some of the GUI buttons are incorrectly enabled or disabled.

Workaround: Click the cluster's Refresh button.

Refresh Issues in Pages Related to Zone Clusters (19144804)

Problem Summary: Sometimes the GUI does not refresh the data in the zone cluster pages. As a result, some of the GUI buttons are incorrectly enabled or disabled.

Workaround: Restart the Oracle GlassFish Server using the following command:

```
# svcadm restart svc:/system/cluster/manager-glassfish3:default
```

Resource Creation Using the Resource Wizard in the GUI Sometimes Fails With a NullPointerException (19130707)

Problem Summary: Sometimes the resource wizard in the GUI fails to configure a resource and fails with a NullPointerException. The following error message is displayed:

```
[AccessBase.getMBean() failed due to exception: java.lang.NullPointerException]
```

At this point, you will not be able to use the resource wizard to create a new resource.

Workaround: Log out of the GUI and log back in.

At Installation, cluster/manager:default Fails Fatally and Transitions to Maintenance Mode (18616858)

Problem Summary: After you install the Oracle Solaris Cluster software, the cluster/manager:default SMF service might end up in the maintenance mode.

Workaround: For the time period between the Oracle Solaris Cluster software installation and the initial cluster configuration, you can ignore this problem.

Although the cluster/manager:default SMF service deploys the Oracle Solaris Cluster Manager web-application (GUI), no services are provided by the GUI until the initial cluster configuration (using /usr/cluster/bin/scinstall). Initial cluster configuration will reboot the node and the SMF service will come online properly at that time.

Property Editing: New Values Take One Refresh Cycle to Appear After Being Saved (18791809)

Problem Summary: After clicking Save to save property edits, the display is updated but contains the previous data.

Workaround: Wait for the next refresh cycle (10 seconds or less) and the new data will appear.

Unable to Install Oracle Solaris Cluster for solaris10 Branded Zone (19064831)

Problem Summary: If a solaris10 branded zone cluster is installed with an archive that does not contain the Oracle Solaris Cluster software, you cannot install the software separately using the GUI.

Workaround: Use the `clzonecluster install-cluster` command to install the Oracle Solaris Cluster software on a solaris10 brand of zone cluster.

No Resource Types Listed in the Pull-Down Menu When Creating a New Resource Through the GUI Wizard (19013853)

Problem Summary: When the resource creation wizard is executed within a zone cluster, the menu to select the resource type can be empty.

Workaround: Restart the Oracle GlassFish Server using the following command:

```
# svcadm restart svc:/system/cluster/manager-glassfish3:default
```

Resource Group Not Listed in the GUI When It Is in the Protection Group But Has Been Removed From the Cluster (18720984)

Problem Summary: If the application resource group in the protection group is removed from the cluster, the protection group experiences a configuration error. To fix this problem, you can either re-create the resource group or remove the resource group from the protection group. Removing the resource group from the protection group is not possible with the GUI because it does not list the resource group in the protection group page and the Remove button is grayed out.

Workaround: Use the `geopg remove-resource-group resource-group-list protection-group-name` command to remove the resource group from the protection group. This command will bring the protection group into the correct configuration.

Very Long Wait Times When Navigating to Partnership and Protection Group If a Cluster in the Partnership Is Down (18961504)

Problem Summary: Navigating to the partnership page or adding a replication component takes three to five minutes when the partner cluster is down.

Workaround 1: Do not administer or manage a partnership whose partner cluster is down. Otherwise, wait for the time out.

Workaround 2: Use the CLI to administer or manage a partnership whose partner cluster is down.

Data Services Notes

This chapter contains known issues and bugs that affect data services of Oracle Solaris Cluster 4.2. Contact Oracle support services to find out whether a code fix has become available.

Data Services Issues

Oracle Database/WLS Resource Fails to Come Online Due to Locking Issue (15713853)

Problem Summary: When using a ZFS Storage Appliance, during a power failure test, after powering off all the cluster nodes and then powering them back on, the database might not come back online and the whole application might fail. Whenever a power cycle happens, the application might not be available until you manually clear the NFS locks from ZFS Storage Appliance storage.

Workaround: For ZFS Storage Appliance storage (NFS file systems), from the ZFS Storage Appliance GUI, go to maintenance, select workflows and then click Clear Locks (with hostname and IP address).

Disabled LDom Resource Is Keeping the LDom in Bound State (16323723)

Problem Summary: This issue does not affect switchovers or failovers but does affect the manual maintenance operations. When a logical domain (LDom) resource is disabled and the zpool that is hosting it needs to be exported, the zpool export fails because the bound state of the LDom keeps the zpool (dependent HAStoragePlus resource) in a busy state.

Workaround: Perform the following steps.

1. Unbind the LDom.

```
# ldm unbind-dom
```

2. Clear the HASStoragePlus resource STOP_FAILED error.

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

Register Scripts Need to Define resource_offline_restart Dependency (15700470)

Problem Summary: The register script defines a resource_dependency to the HASStoragePlus resource but This dependency needs to be a resource_dependency_offline_restart dependency.

Workaround: Disable the HA for the PostgreSQL resource and remove the HASStoragePlus resource from the list of resource_dependencies. Then, add the HASStoragePlus resource to the resource_dependencies_offline_restart property.

NFS Resource and HASP Resources Fail in Oracle Solaris Cluster 4.2 on Oracle Solaris 11.1 and Oracle Solaris 11.2 With zfs recv (17365301)

Problem Summary: This issue might occur on a system configured with a SUNW.HASStoragePlus (HASP) resource managing a ZFS storage pool (zpool) and a SUNW.nfs resource managing an NFS share of a ZFS file system from this zpool. For more information, see the [SUNW.nfs\(5\)](#) man page.

When a large zfs send and zfs recv is performed with a snapshot from another system to a separate ZFS sub-volume on the same zpool that is managed by HASP, both the NFS and the HASP resources might fail in Oracle Solaris Cluster 4.2 running on Oracle Solaris 11.1 or Oracle Solaris 11.2.

Workaround: Before starting data replication of the file system that is actively managed under the Oracle Solaris Cluster resource, do either of the following:

- Execute the following command to disable the NFS resource:

```
# clresource disable nfs-resource-name
```

- Execute the following command to disable monitoring of the NFS resource:

```
# clresource unmonitor nfs-resource-name
```

Once data replication is successfully completed, bring the NFS resource to a monitored and online state.

Data Service Configuration Wizards Do Not Support Storage Resources and Resource Groups for Scalable HAStoragePlus (15820415)

Problem Summary: The existing data service configuration wizards do not support configuring scalable HAStoragePlus resources and resource groups. In addition, the wizards are also not able to detect existing resources and resource groups for scalable HAStoragePlus.

For example, while configuring HA for WebLogic Server in multi-instance mode, the wizard will display `No highly available storage resources are available for selection` even when there are existing scalable HAStoragePlus resources and resource groups on the cluster.

Workaround: Configure data services that use scalable HAStoragePlus resources and resource groups as follows:

1. Use the `clresourcegroup` and `clresource` commands to configure HAStoragePlus resources groups and resources in scalable mode.
2. Use the `clsetup` wizard to configure data services as if they are on local file systems, meaning as if no storage resources are involved.
3. Use the CLI to create an offline-restart dependency on the scalable HAStoragePlus resources configured in Step 1, and a strong positive affinity on the scalable HAStoragePlus resource groups.

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` Command on Multiple Nodes at the Same Time Could Cause Problem (15533979)

Problem Summary: 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 any one active cluster node.

Developer Environment Notes

This chapter contains known issues and bugs that affect the developer environment of Oracle Solaris Cluster 4.2. Contact Oracle support services to find out whether a code fix has become available.

Developer Environment Issues

num_zoneclusters Property Cannot Be Set in Non-cluster Mode (1852819)

Problem Summary: Currently, the num_zoneclusters property can be set only in the cluster mode and not in the non-cluster mode.

Workaround: If the num_zoneclusters property needs to be set or changed, boot the node in cluster mode.

Geographic Edition Notes

This chapter contains known issues and bugs affect the Oracle Solaris Cluster Geographic Edition 4.2 software. Contact Oracle support services to find out whether a code fix has become available.

Geographic Edition Issues

geosite create Fails With JMX timeout, Dangling Heartbeat is Left Behind After geosite delete (19232931)

Problem Summary: When the `geosite create` command fails due to the JMX timeouts, the site is not cleaned up.

Workaround: Perform the following steps:

1. If the site still exists on the cluster, run the following command:

```
# geosite leave -f site-name
```

2. Restart common agent container on the node to cleanup stale heartbeats where the failover control resource is online.

```
# cacoadm restart
```

3. Add trust manually for each cluster that you want to configure on the site.

```
# geoops add-trust -c cluster-name
```

4. Re-create the site.

```
# geosite create site-name
```

Switchover Proceeds and Start Reports Success When the SRDF Device Group Is in a Partitioned State (17250376)

Problem Summary: If there is a link failure between the EMC storage boxes, then the SRDF device group is in a partitioned state. In this situation, running the `geopg switchover` command is successful. However, the command to swap the SRDF device group fails.

Workaround: Run the `symrdf failover` and the `symrdf swap` commands so that the device group, `devgroup2`, has the RDF2 role on the new primary site.

```
# symrdf -g devgroup2 failover
# symrdf -g devgroup2 swap
```

Oracle Data Guard Module Incorrectly Flagging SUNW.oracle_server Dependencies in the Single Instance (15818725)

Problem Summary: Attempting to retrieve the Oracle Data Guard protection group configuration fails with an error if HA for Oracle Database has dependencies on other resources.

Workaround: Set the protection group's `external_dependencies_allowed` property to `true`.

```
# geopg set-prop -p external_dependencies_allowed=TRUE protection_group
```

Data Guard Module Incorrectly Updates rg_affinities and Protection Group Validation Failures When Multiple rac_server_proxy Resource Groups Exist on the Cluster (18940623)

Problem Summary: Updating the Oracle Data Guard replication component to a new `rac_server_proxy` resource group in the protection group appends with the old `rac_server_proxy` and results in `geopg validate protection-group` failures.

Workaround: Set the protection group's `external_dependencies_allowed` property to `true`.

```
# geopg set-prop -p external_dependencies_allowed=TRUE protection_group
```


Geographic Edition Runs show configuration Frequently and Results in the Database Getting Requests Seconds Apart (18779118)

Problem Summary: Frequent show configuration requests from the Oracle Data Guard replication resource probe method leads to ORA-00020 errors and delayed takeover.

Workaround: To avoid ORA-00020 errors, ensure that you have applied Oracle Database patch 18746642. Also, increase the parameter processes by running the following command:

```
SQL> alter system set processes=new_max;
```

Collision Problems Should Be Flagged at Protection Group Creation Time (15801862)

Problem Summary: Projects or mount points configured with the same name on the target appliance as the one on source appliance managed by Geographic Edition on the primary cluster will result in switchover or takeover failures.

Workaround: Before adding the Oracle ZFS Storage Appliance replicated project to the protection group, ensure that the target appliance does not have a project or mount point with the same name as the source appliance.

If Operation on Any Protection Group Fails During a Multigroup Operation, Then the Multigroup Operation Errors Out Without Doing Any Pending Operations on Other Protection Groups in the Multigroup (18270451)

Problem Summary: A multigroup operation (start, stop, switchover, and takeover) performs required operations on the protection groups in the multigroup. The multigroup performs such operations in parallel or sequentially depending on the type of operation and the dependencies between the protection groups as specified in the multigroup. If any such protection group operation fails, then the multigroup operation errors out without performing any pending operations. A multigroup operation prints messages indicating when an operation on a protection group starts and when it ends. It also reports any operation failure to the user.

Workaround: Besides the failure messages reported by multigroup operation, check the status of the individual protection groups to see whether the requested operation has been performed.

Fix any issues that caused the operation to fail on a protection group. Then, use the `geopg` command to perform individual protection group operation on such protection groups.

If Resource Group Is in Pending_online_blocked State, geopg start Continuously Tries to Start the Resource Group (16636516)

Problem Summary: When the `geopg start` command is run and a resource group contained in the protection group is in `Pending_online_blocked` state, then the operation fails with an attempt to start the resource groups whose dependent resource groups are in the `Online Failed/starting` state.

Workaround: Ensure that the dependent resource groups of the resource groups in the protection groups are in `Online` or `Offline` state before running the `geopg start` command.

After Changing Timeout of a Protection Group, an Attempt to Change the Timeout of a Multigroup Containing That Protection Group Fails in Some Situations (18436909)

Problem Summary: The timeout of a multigroup is based on the timeouts of the protection groups in the multigroup. If the timeout of a protection group has been changed but the configuration of a multigroup that contains the protection group has not been refreshed (through a `validate` command, for instance), then an attempt to set the timeout of the multigroup errors out based on the older timeout of the protection group as known to the multigroup. This situation happens when the new proposed multigroup timeout does not conform with the earlier known timeouts of the constituent protection groups.

Workaround: Refresh the configuration of the multigroup using the `geomg validate` command and then set the proposed timeout on the multigroup using the `geomg set-prop` command.

Disallow `site create` or `site add-member` When Site of This Name Already Exists on Members (18415502)

Problem Summary: Allowing the issuing of the `site create` or `site add-member` command when the controller or member being added to the configuration already has a site of this name results in the site being rejected by the clusters (controllers or members) when an attempt is made to synchronize and update.

Workaround: Ensure that a site with same name is not created on any of the controllers or members that are added to the configuration.

Doing `geosite update remote-cluster site` on a Cluster Does Not Replicate the site's Multigroups That Are Present on the Remote Cluster onto the Local Cluster (18368896)

Problem Summary: Once a multigroup is created using `geomg create` on any controller in a site, the multigroup gets created automatically on other clusters in the site if that controller has no site configuration synchronization errors with those clusters. If the site synchronization status is in ERROR between any such cluster and that controller, then that cluster does not accept the multigroup creation.

One possible way to attempt to resolve the site synchronization ERROR is by using the `geosite update` command on that cluster with the controller as an argument in order to make the site's configuration data on the cluster the same as the data that exists on the controller, and thereby replicate the multigroup onto that cluster. This replication of a multigroup configuration might fail in some situations even though the site synchronization status of that cluster will report OK with respect to the controller.

Workaround: Use the `geosite leave` command to make that cluster leave the site and then include it back in the site using the `geosite add-member` and `geosite join` commands.

Documentation Notes

This chapter discusses errors or omissions for documentation in the Oracle Solaris Cluster 4.2 release.

System Administration Guide

Chapter 11, “Updating Your Software,” in “Oracle Solaris Cluster System Administration Guide” is missing instructions for updating a zone cluster that is not a `solaris` or `solaris10` brand of zone cluster. To update a zone cluster that is not branded, follow procedures in “Updating a Specific Package” in “Oracle Solaris Cluster System Administration Guide” to update the underlying global cluster. When the global cluster is updated, its zone clusters are automatically updated as well.

HA for SAP NetWeaver Guide

- In “Registering and Configuring the HA for SAP NetWeaver by using Oracle Solaris Cluster Command Line Interface (CLI)” in “Oracle Solaris Cluster Data Service for SAP NetWeaver Guide”:
 - SAP dialogue instance must not be included as one of the SAP central services components.
 - `additional dialogue instance` must read as `additional application server instance`.
- In “How to Register and Configure an SAP NetWeaver `sapstartsrv` Resource” in “Oracle Solaris Cluster Data Service for SAP NetWeaver Guide”:
 - Step 2, replace `-p instance_number=instance \` with `-p instance_number=instance_number \`.
 - Step 2, replace `-p instance_name=instance_number \` with `-p instance_name=instance_name \`.
 - Step 2, add `-p child_mon_level=5 \` to the existing list of parameters, while creating a failover `sapstartsrv` resource as well as while creating a multiple master `sapstartsrv` resource if you do not configure it for a resource of the resource type `ORCL.sapdia`, which is the primary application server and additional application server resources. If

you configure it for either the primary application server or additional application server resources, then use the following parameter:

`-p child_mon_level=6 \`

- In [“How to Register and Configure an SAP Central Service Resource”](#) in [“Oracle Solaris Cluster Data Service for SAP NetWeaver Guide”](#):
 - Step 3, replace `-p resource_dependencies=db-rs,scs-startsrv-rs \` with `-p resource_dependencies=scs-startsrv-rs \`.
 - Step 3, add `-p retry_count=0 \` to the existing list of parameters. This parameter specifies that the central services resource must not be restarted in case of a failure and must perform a failover instead.
 - Step 3, replace `-p instance_number=instance \` with `-p instance_number=instance_number \`.
 - Step 3, replace `-p instance_name=instance_number \` with `-p instance_name=instance_name \`.
- In [“How to Register and Configure an SAP Replicated Enqueue Server Resource”](#) in [“Oracle Solaris Cluster Data Service for SAP NetWeaver Guide”](#):
 - Step 3, replace `-p resource_dependencies=db-rs,rep-startsrv-rs \` with `-p resource_dependencies=rep-startsrv-rs \`.
- In [“How to Register and Configure an SAP Replicated Enqueue Preempter Resource”](#) in [“Oracle Solaris Cluster Data Service for SAP NetWeaver Guide”](#):
 - Step 2, replace `-p enq_instnr=instance_list` with `-p enq_instnr=instance_number_list \`.
 - Step 2, remove `-p resource_dependencies=db-rs \` from the list of parameters.
 - Step 2, replace `rep-rs` with `preempt-rs`.
- In [“How to Register and Configure an SAP NetWeaver Application Server Instance Resource”](#) in [“Oracle Solaris Cluster Data Service for SAP NetWeaver Guide”](#):
 - The title of this section should read "How to Register and Configure an SAP NetWeaver Application Server Instance Resource."
 - `primary application instance` should read `primary application server instance`.
 - `additional application instance` should read `additional application server instance`.
 - Instances containing dialogue instances and central services can be combined as central instance without the central services resources. You must add the following property:

`-p Architecture=comb \`
- In [“How to Enable Interaction Between Oracle Solaris Cluster and SAP NetWeaver”](#) in [“Oracle Solaris Cluster Data Service for SAP NetWeaver Guide”](#):
 - Step 1 should read as follows:

1) Configure every instance profile with the following lines, assuming that the script is in the `/usr/sap/SID/SYS/exe/run` directory.

```

#
# SAP HA Script Connector
#
service/halib = /usr/sap/<SID>/SYS/exe/run \
service/halib_cluster_connector = \
/opt/ORCLscsapnetw/saphacmd/bin/sap_orcl_cluster_connector \
service/halib_debug_level = 1

```

Man Pages

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

cldevicegroup(1CL)

The `cldevicegroup` man page is updated with the following information for Oracle Solaris Cluster 4.2 SRU 1:

<code>create</code>	<p>Creates a new device group.</p> <p>You can use this subcommand only in the global zone.</p> <p>This subcommand supports only the <code>rawdisk</code> device-group type. You cannot create an <code>svm</code> or <code>sds</code> device group by using Oracle Solaris Cluster commands. Instead, use Solaris Volume Manager commands to create Solaris Volume Manager disk sets. Disk sets are automatically registered with Oracle Solaris Cluster software as <code>svm</code> or <code>sds</code> device groups. For more information about device-group types, see the description of the <code>-t</code> option.</p> <p>If you specify a configuration file with the <code>-i</code> option, you can supply a plus sign (+) as the operand. When you use this operand, the command creates all device groups that are specified in the configuration file that do not already exist.</p> <p>For device groups of type <code>rawdisk</code>, use the <code>-d</code> option with the <code>create</code> subcommand to specify one or more devices to the device group. You cannot create a device group without any device. When you specify devices, use one <code>-d</code> option per command invocation. You cannot create multiple raw-disk device groups in one command invocation unless you use the <code>-i</code> option.</p> <p>Users other than superuser require <code>solaris.cluster.modify</code> RBAC authorization to use this subcommand.</p>
---------------------	--

For information about how to delete device groups, see the description of the `delete` subcommand.

`remove-device`

Removes member disk devices from a raw-disk device group.

You can use this subcommand only in the global zone.

The `remove-device` subcommand is only valid with device groups of type `rawdisk`. This subcommand is not valid with `svm` or `sds` device-group types.

You cannot use the `remove-device` subcommand to remove all the devices in a device group. A device group must contain at least one device. To remove all the devices in a device group, use the `delete` subcommand to remove the device group.

Users other than superuser require `solaris.cluster.modify` RBAC authorization to use this subcommand.

For information about how to add disk devices to a raw-disk device groups, see the description of the `add-device` subcommand.

The value of the `num_xip_zoneclusters` property cannot be less than the highest assigned `clprivnet` instance number.

cluster(1CL)

The following information is omitted from the `num_xip_zoneclusters` property description:

The value of the `num_xip_zoneclusters` property cannot be less than the highest assigned `clprivnet` instance number.