

Oracle® Solaris Cluster 4.3 Release Notes

ORACLE®

Part No: E56685
June 2021

Part No: E56685

Copyright © 2000, 2021, Oracle and/or its affiliates.

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, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. 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 Inside 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, Epyc, and the AMD 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 about 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 unless otherwise set forth in an applicable agreement between you and Oracle. 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, except as set forth in an applicable agreement between you and Oracle.

Documentation Accessibility

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

Access to Oracle Support

Oracle customers that have purchased support 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.

Diversity and Inclusion

Oracle is fully committed to diversity and inclusion. Oracle recognizes the influence of ethnic and cultural values and is working to remove language from our products and documentation that might be considered insensitive. While doing so, we are also mindful of the necessity to maintain compatibility with our customers' existing technologies and the need to ensure continuity of service as Oracle's offerings and industry standards evolve. Because of these technical constraints, our effort to remove insensitive terms is an ongoing, long-term process.

Référence: E56685

Copyright © 2000, 2021, Oracle et/ou ses affiliés.

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 stipulation expresse de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, accorder de licence, 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 livré sous licence au Gouvernement des Etats-Unis, ou à quiconque qui aurait souscrit la licence de ce logiciel 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 embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. 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 un risque de 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 des applications dangereuses.

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 Inside 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, Epyc, et le logo AMD sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée de 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, sauf mention contraire stipulée dans un contrat entre vous et Oracle. 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, sauf mention contraire stipulée dans un contrat entre vous et Oracle.

Accessibilité de la documentation

Pour plus d'informations sur l'engagement d'Oracle pour l'accessibilité de la documentation, visitez le site Web Oracle Accessibility Program, à l'adresse : <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Accès aux services de support Oracle

Les clients Oracle qui ont souscrit un contrat de support ont accès au support électronique via My Oracle Support. Pour plus d'informations, visitez le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> ou le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> si vous êtes malentendant.

Contents

Using This Documentation	9
Oracle Solaris Cluster 4.3 Notes	11
What's New in the Software	11
Former Features Not Included in the Oracle Solaris Cluster 4.3 Software	17
Product Localization Notes	17
Commands Modified in This Release	18
Compatibility Notes	21
Compatibility Issues	21
Restriction of Support for StorageTek QFS	21
Restriction of the Use of EMC SRDF in a Campus Cluster	21
FIPS-140 Compatibility	21
The Output of the <code>verifyrwprojects</code> API Was Changed From the Original Format (23108669)	22
Restriction of Oracle Grid Infrastructure Support for DLMP Link Aggregations or VNICs in Shared-IP Zones (21660315)	22
<code>md_stripe: WARNING: md: write error, md: Panic due to lack of DiskSuite state (21785654)</code>	23
<code>listrwprojects</code> Client Interface Not Listing Projects That Have Extra IPs (19982694)	23
VNIC Names Longer Than 16 Characters Cause Problems (17362337)	23
Zone Does Not Boot if <code>pkg:/system/resource-mgmt/resource-cap</code> Is Not Installed and <code>capped-memory</code> Is Configured (15740089)	24
Active:Active ZFS Storage Appliance Clustered Configurations Are Not Supported (15521899)	24
Installation Notes	25

Installation Issues	25
Same DID Device Being Created for All Boot Disks When Installed With Oracle VM Server for SPARC Template (23755653)	25
scinstall Fails While Using Only IPv6 for a Public Network Configuration (16355496)	26
Hard to Determine Data Service Names for solaris10 Branded Zone Noninteractive Data Service Installation (15804349)	26
Administration Notes	29
Administration Issues	29
Unable To Reset acfs_mountpoint Property When the SUNW. scalable_acfs_proxy Resource Is Disabled (16928490)	29
clzonecluster install -a archive-no-cluster-pkgs zone-cluster Does Not Install Cluster Packages (18714803)	29
IPv6 Scalable Service Support Is Not Enabled By Default (15290321)	30
Removing a Node From an Exclusive-IP Zone Cluster Panics the Cluster Nodes (15817184)	30
EMC SRDF Rejects Switchover When Replicated Device-Group Status Will Cause Switchover and Switchback to Fail (15538295)	31
clzonecluster apply Fails to Add Device and Filesystem Into the shared-ip Zone with Errors (21541048)	31
Runtime Notes	33
Runtime Issues	33
clcomm: path online, Then Path Being Drained Every Eleven Seconds (18827672)	33
Oracle Solaris Cluster Manager Browser Interface Notes	35
Browser Interface Issues	35
HA-Zones Wizard Should Permit Kernel Zones Live Migration for SPARC (23025005)	35
Unable to Install Oracle Solaris Cluster for solaris10 Branded Zone (19064831)	35
Oracle Solaris Cluster Manager Configuration Wizard Configures Only Three Nodes Out of the Four (21490228)	36
Oracle Solaris Cluster Manager Browser Interface Cannot Run Under Trusted Extensions (21323252)	36

Oracle Solaris Cluster Manager Browser Interface Reports a Connection to the Server Failed Error (21480830)	36
Deleting Zone Cluster by Using <code>clzonecluster</code> Causes the Oracle Solaris Cluster Manager Page to Loop (21555137)	37
Errors DCA-29000, http 500, JBO-29114 Displayed If User Removes the Current Cluster From Site on Site Protection Group (21661908)	37
Data Services Notes	39
Data Services Issues	39
Oracle Database/WLS Resource Fails to Come Online Due to Locking Issue (15713853)	39
HASP Resources Fail in Oracle Solaris Cluster 4.3 on Oracle Solaris 11.2 and Oracle Solaris 11.3 With <code>zfs recv</code> (17365301)	39
Data Service Configuration Wizards Do Not Support Storage Resources and Resource Groups for Scalable HAStoragePlus (15820415)	40
Scalable Applications Are Not Isolated Between Zone Clusters (15611122)	41
NFS Server Failover Triggers Stale NFS File Handle (21459179)	41
Upgrading From Oracle Solaris 11.2 to Oracle Solaris 11.3 Results in Oracle Grid 12.1.0.1.0 Startup Hang (21511528)	41
ORA-00742: Log Read Detects Lost Write (21186724)	42
Developer Environment Notes	43
Developer Environment Issues	43
<code>num_zoneclusters</code> Property Cannot Be Set in Non-cluster Mode (18528191)	43
Geographic Edition Notes	45
Geographic Edition Issues	45
Case Where DIDs Get Reassigned in a Geographic Edition Oracle ZFS Storage Appliance Configuration (24851015)	45
Oracle Data Guard Module Incorrectly Flagging <code>SUNW.oracle_server</code> Dependencies in the Single Instance (15818725)	46
Collision Problems Should Be Flagged at Protection Group Creation Time (15801862)	47
Doing <code>geosite update remote-cluster site</code> on a Cluster Does Not Replicate the Site's Multigroups That Are Present on the Remote Cluster Onto the Local Cluster (18368896)	47

Infrastructure Resource Goes Offline After Probe Caused a Restart and Failed to Start (21298474)	48
Oracle GoldenGate Protection Group Data Replication Status Shows OK Though the Replication Resource did not Failover (21527062)	48
java.lang.IllegalArgumentException: Unmatched braces in the pattern (21570583)	49
Protection Group Creation Should Not Fail If a Node in the Cluster Is Down (21697993)	49
If Takeover Is Performed While Both Sites Are Up, Project Is Not Removed From the Original Primary Site (21684495)	49
Geographic Edition Does Not Support ZFSSA Offline Replication Feature (21934145)	50
Documentation Notes	51
Man Pages	51
Extraneous Characters in the Command Lines of Some Man Page Examples	51
geopg(1M) Man Page	51
New and Enhanced Extension Properties for the SUNW.Proxy_SMF_* Resource Types	52
Data Services Planning and Administration Guide	54
HA for Oracle External Proxy Guide	55
HA for PostgreSQL Guide	56
Hardware Administration Manual	56
Software Installation Guide	56
System Administration Guide	56
Addendum to Support Network Appliance NAS Devices	56

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

Documentation and resources for this product and related products are available at <http://www.oracle.com/pls/topic/lookup?ctx=E56676-01>.

Feedback

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

Diversity and Inclusion

Oracle is fully committed to diversity and inclusion. Oracle respects and values having a diverse workforce that increases thought leadership and innovation. As part of our initiative to build a more inclusive culture that positively impacts our employees, customers and partners we are working to remove insensitive terms from our products and documentation. We are also mindful of the necessity to maintain compatibility with our customers' existing technologies and the need to ensure continuity of service as Oracle's offerings and industry standards evolve.

Because of these technical constraints, our effort to remove insensitive terms is ongoing and will take time and external cooperation.

Oracle Solaris Cluster 4.3 Notes

This document provides the following information about the Oracle Solaris Cluster 4.3 software:

- “What's New in the Software” on page 11
- “Former Features Not Included in the Oracle Solaris Cluster 4.3 Software” on page 17
- “Product Localization Notes” on page 17
- “Commands Modified in This Release” on page 18
- “Compatibility Issues” on page 21
- “Installation Issues” on page 25
- “Administration Issues” on page 29
- “Runtime Issues” on page 33
- “Browser Interface Issues” on page 35
- “Data Services Issues” on page 39
- “Geographic Edition Issues” on page 45
- “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 4.3 software. For other changes, also see “[Commands Modified in This Release](#)” on page 18 and “[Documentation Notes](#)”.

The Oracle Solaris Cluster 4.3 software provides the following new features:

- **IPSec support for private interconnects** – Beginning with Oracle Solaris Cluster 4.3.7, you can secure cluster private interconnect traffic by configuring IPSec. IPSec configuration protects the inter-node traffic for internal cluster operations as well as user data for global

filesystems. For more information about configuring IPSec, see [“Securing the Interconnect for Oracle Solaris Cluster with IPSec”](#) in *Oracle Solaris Cluster 4.3 Software Installation Guide*.

- **Easy updation of ZFS Storage Appliance Data Replication properties** – Beginning with Oracle Solaris Cluster 4.3.7, you can use registration script to update data replication component properties. An additional parameter `MODIFY_PASSPHRASE` is used in ZFSSA configuration file which allows you to modify the password during an attempt to modify the existing replication component.
- **Geographic Edition Clusters now monitor ZFS Storage Appliance Cluster status in probe and switchover operations** – Beginning with Oracle Solaris Cluster 4.3.7, Geographic Edition Cluster probe and pre-checks for switchover monitor the status of ZFSSA cluster heads. They fail if ZFSSA cluster heads are not in clustered state.
- **Support for ZOSS (Zones on Shared Storage) over NFS from ZFS Storage Appliance in the HA Zone GUI wizard for kernel zones** – Beginning with Oracle Solaris Cluster 4.3.7, you can configure Oracle Solaris kernel zones whose zonpath resides over NFS as highly available. This feature enables you to create and manage the required resources for the corresponding NFS storage associated with the zone.
- **Support for setting ZFS Storage Appliance NFS exceptions in the NAS GUI wizard** – Beginning with Oracle Solaris Cluster 4.3.7, when using the NAS GUI wizard to add projects to the cluster, the wizard automatically updates the project's NFS share mode and NFS exception list to meet cluster requirements.
- **Example Files for Customizing an Automated Installer Installation of a Cluster** – Beginning with Oracle Solaris Cluster 4.3.6, you can create or modify a custom Automated Installer (AI) manifest and system configuration profile for cluster installation. The custom files can be used as an alternative to an AI installation using the `scinstall` utility.
See [Appendix A, “Example Automated Installer Cluster Installation Files,”](#) in *Oracle Solaris Cluster 4.3 Software Installation Guide* for guidelines and examples.
- **Oracle Solaris Cluster Manager support of user rights profiles** – Beginning with Oracle Solaris Cluster 4.3.6, the Oracle Solaris Cluster Manager browser interface supports users or roles that are assigned rights profiles for functionality performed in the interface. The interface also restricts certain Oracle Solaris Cluster functions from users or roles that do not have sufficient assigned rights profiles to perform those functions. For more information, see [“User Rights for Oracle Solaris Cluster Manager”](#) in *Oracle Solaris Cluster 4.3 System Administration Guide*.
- **Restriction change for the subnet used by a ZFS Storage Appliance NAS device** – An additional option for the subnet location of a ZFS Storage Appliance NAS device or quorum device is qualified for the 4.3 release. Support is added for the device to be located on a different subnet than the cluster nodes' subnet, provided that L3 switching is used between the subnets. This change is documented in [“Requirements When Configuring Oracle ZFS Storage Appliances”](#) in *Oracle Solaris Cluster With Network-Attached Storage Device Manual* and [“Requirements When Configuring Oracle ZFS Storage Appliance NAS](#)

Devices as Quorum Devices” in *Oracle Solaris Cluster With Network-Attached Storage Device Manual*.

- **New extension properties for the SUNW.Proxy_SMF_* resource types** – Beginning with Oracle Solaris Cluster 4.3.4 software, the SUNW.Proxy_SMF_failover, SUNW.Proxy_SMF_multimaster, and SUNW.Proxy_SMF_scalable resource types support the following new extension properties:
 - Monitor_retry_count
 - Monitor_retry_interval
 - Probe_command

In addition, the Probe_timeout extension property can now be specified directly to the clresource command.

For more information about these properties, see “[New and Enhanced Extension Properties for the SUNW.Proxy_SMF_* Resource Types](#)” on page 52.

- **Oracle VM Server for SPARC template to configure a guest domain or I/O domain cluster** – Beginning with Oracle Solaris Cluster 4.3.4 software, a pre-built Oracle VM Server for SPARC template for Oracle Solaris Cluster is available. to configure a new cluster of guest domains or I/O domains. The template includes software packages for Oracle Solaris 11.3 and Oracle Solaris Cluster 4.3. The template can be downloaded and deployed to both create the new logical domains and to install and configure those domains to form a cluster.

For information about downloading and using the template, see “[Establishing a New Logical Domain Cluster by Deploying the Oracle Solaris Cluster Oracle VM Server for SPARC Template](#)” in *Oracle Solaris Cluster 4.3 Software Installation Guide*.

- **Additional enhancements to the Oracle Solaris Cluster Manager Browser Interface** – Beginning with Oracle Solaris Cluster 4.3.4, the following functionality is added to Oracle Solaris Cluster Manager:
 - New configuration wizard for the Oracle WebLogic Server data service
 - Support of Oracle Solaris Cluster Geographic Edition in zone clusters
 - Support direct authentication into a zone cluster

For more information about Oracle Solaris Cluster Manager, see *Oracle Solaris Cluster 4.3 System Administration Guide*.

- **Support for Oracle Data Guard Far Sync Instances With Oracle Solaris Cluster Geographic Edition** – Beginning with Oracle Solaris Cluster 4.3.4 software, an Oracle Data Guard far sync instance is supported for data replication in an Oracle Solaris Cluster Geographic Edition disaster recovery partnership.

For more information about using Oracle Data Guard data replication in a disaster recovery partnership, see *Oracle Solaris Cluster Geographic Edition Data Replication Guide for Oracle Data Guard*. For information about creating an Oracle Data Guard far sync instance

and the supported protection modes for far sync instances, see [“Far Sync” in Oracle Data Guard Concepts and Administration](#).

- **Support for Hitachi TrueCopy with Oracle Solaris Cluster Geographic Edition** – Beginning with Oracle Solaris Cluster 4.3.2 software, Hitachi TrueCopy data replication is supported in a Geographic Edition configuration. For more information, see [Oracle Solaris Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy and Universal Replicator](#).
- **Enhancements to the Oracle Solaris Cluster Manager Browser Interface** – Beginning with Oracle Solaris Cluster 4.3.3 software, the following functionality is added to Oracle Solaris Cluster Manager:
 - New configuration wizard for the Support of Oracle RAC data service
 - New configuration wizard for per-node logical hostnames
 - New configuration wizard for storage resources for a zone cluster
 - Support for setting kernel zones live migration in the HA for Oracle Solaris Zones configuration wizard (initial support for x86 clusters only)

For more information about Oracle Solaris Cluster Manager, see [Oracle Solaris Cluster 4.3 System Administration Guide](#).

- **Support for Oracle Solaris Cluster Data Service for Oracle E-Business Suite 12.2** – Beginning with Oracle Solaris Cluster 4.3.3 software, Oracle E-Business Suite 12.2 release is supported.

For more information, see [Oracle Solaris Cluster Data Service for Oracle E-Business Suite as of Release 12.2 Guide](#).

For Oracle E-Business Suite versions up to 12.1, see [Oracle Solaris Cluster Data Service for Oracle E-Business Suite up to Release 12.1 Guide](#).

- **Support for Oracle Solaris Cluster Data Service for Oracle Siebel 8.1.1.11 and 8.1.1.14** – Beginning with Oracle Solaris Cluster 4.3.3 software, Oracle Siebel 8.1.1.11 and 8.1.1.14 are supported.

For more information, see [Oracle Solaris Cluster Data Service for Siebel Guide SPARC Platform Edition](#).

- **Support for EMC Symmetrix Remote Data Facility in zone cluster configuration** – Beginning with Oracle Solaris Cluster 4.3.3 software, EMC SRDF is supported in the zone cluster configuration.

For more information, see [Oracle Solaris Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility](#).

- **Qualification of Hitachi Universal Replicator in a Geographic Edition Configuration** – Beginning with Oracle Solaris Cluster 4.3.2 software, Hitachi Universal Replication is qualified to use for storage-based data replication in a Geographic Edition configuration. As of the 4.3.2 release, the following restrictions apply:
 - Geographic Edition with Hitachi Universal Replicator using HA-NFS is not supported.

- Hitachi Universal Replicator use in a campus cluster is not qualified.
- Use of Hitachi TrueCopy in campus cluster or Geographic Edition is not qualified.

For more information, see [Oracle Solaris Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy and Universal Replicator](#).

- **Support for Oracle MaxRep for SAN replication in Oracle Solaris Cluster Geographic Edition** – Beginning with Oracle Solaris Cluster 4.3.3 software, Oracle Solaris Cluster Geographic Edition supports data replication by using Oracle MaxRep for SAN in a disaster recovery setup. You can set up Geographic Edition protection groups with MaxRep replication for protection plans to protect application data.

For more information, see [Oracle Solaris Cluster Geographic Edition Data Replication Guide for Oracle MaxRep for SAN](#).

- **Enhanced Oracle Solaris Cluster Manager** – In this release, Oracle Solaris Cluster Manager is capable of performing following actions:
 - The 4.3 release of Oracle Solaris Cluster Manager is capable of managing version 4.2 clusters.
 - Oracle Solaris Cluster Manager support for Geographic Edition expanded to include the orchestration feature.
 - New Oracle Solaris Cluster Manager wizard for initial cluster configuration.
 - New Oracle Solaris Cluster Manager wizard for configuring highly available zones.
 - New Oracle Solaris Cluster Manager wizard for configuring highly available Oracle VM Server for SPARC logical domains.
- **Support for importing installed zones as nodes of a new or existing zone cluster** – This feature enables you to make an already installed zone part of a new zone cluster or import the zone to an existing zone cluster configuration.

For more information about zone cluster, see the `clzonecluster(1CL)` man page.

- **Support for Oracle Solaris Zones on shared storage** – In this release, the Oracle Solaris Cluster Data Service for Oracle Solaris Zones supports Oracle Solaris Zones on shared storage. The `sczbt` component now manages non-global zones that specify either the `rootzpool` or `zpool` zone property.

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

- **Support for live migration for kernel zones** – In this release, the Oracle Solaris Cluster Data Service for Oracle Solaris Zones supports live migration for an Oracle Solaris kernel zone on Oracle Solaris 11.3. The `sczbt` component supports live migration for kernel zones if the `Migrationtype` variable in the `sczbt_config` configuration file is set to `live`.

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

- **Support for Oracle Solaris ZFS snapshot data replication in Oracle Solaris Cluster Geographic Edition** – Oracle Solaris Cluster Geographic Edition feature supports data replication by using Oracle Solaris ZFS snapshots in a disaster recovery setup. You can set up Geographic Edition protection groups with ZFS snapshot-based data replication to protect application data.

For more information, see [Oracle Solaris Cluster Geographic Edition Data Replication Guide for ZFS Snapshots](#).

- **Support for Oracle GoldenGate replication in Oracle Solaris Cluster Geographic Edition** – Oracle Solaris Cluster Geographic Edition feature supports data replication by using Oracle GoldenGate.

For more information, see [Oracle Solaris Cluster Geographic Edition Data Replication Guide for Oracle GoldenGate](#).

- **Support for Oracle Solaris Cluster HA for Oracle Communications ASAP** – This high availability data service for Oracle Communications ASAP enables the Oracle Solaris Cluster to start, stop, and monitor ASAP Server components. Oracle Communications ASAP service activation software is used by telecommunications service providers for automated service activation.

For more information, see [Oracle Solaris Cluster Data Service for Oracle Communications ASAP Guide](#) and the `ORCL.asap(5)` man page.

- **Support for Oracle Solaris Cluster HA for Oracle Essbase Server** – This high availability data service for Oracle Essbase Server enables the Oracle Solaris Cluster product to start, stop, and monitor Oracle Essbase Server components. Oracle Essbase Server is a multidimensional database management system used primarily in the financial sector.

For more information, see [Oracle Solaris Cluster Data Service for Oracle Essbase Server Guide](#) and the `ORCL.essbase(5)` man page.

- **Support for timeout threshold notice for Resource Group Manager execution of methods and monitor probes** – Oracle Solaris Cluster enables you to configure timeouts for resource callback methods such as start, stop, or validate, or for monitor probes. A new resource property `Timeout_threshold` enables you to set a timeout threshold representing a percentage of the configured timeout. If a method or probe execution exceeds the configured threshold percentage, an alert is generated, which enables you to adjust the timeout before a failure actually occurs.

For more information about the `Timeout_threshold` property, see the `r_properties(5)` man page.

- **Support for IP over link aggregation for public network** – Oracle Solaris Cluster supports use of regular IP interface over link aggregation as a public network interface. This includes support for IP over DLMP, IP over trunk aggregation, IP over VNIC over DLMP, and IP over VNIC over trunk aggregation. Logical hostname and shared address resources can be assigned any of the supported interface configurations. Logical hostname and shared address wizards automatically discover the supported interfaces that are appropriate for

holding given IP addresses. The `clnode` command displays information about the supported interfaces.

The zone cluster wizard automatically discovers link aggregations and VNICs created over link aggregations that can be assigned to exclusive IP zone cluster.

Former Features Not Included in the Oracle Solaris Cluster 4.3 Software

The following features are included in the Oracle Solaris Cluster 3.3 version but are not included in the Oracle Solaris Cluster 4.3 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 4.3 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	<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>

Commands Modified in This Release

The following commands have been modified in this release.

- Beginning with the Oracle Solaris Cluster 4.3.6 release, the `geopg status` command supports the following additional option:

<code>-v</code>	Displays messages for any errors that put the configuration in the Error state, and messages for any errors that put the replication resource in the FAULTED or DEGRADED error states.
<code>--verbose</code>	
- As of the Oracle Solaris Cluster 4.3.4 release, the `replication_mode` property of the Oracle Data Guard module to Geographic Edition is no longer used and has been removed from the `geopg` command. Ignore mentions of this data replication component property in Oracle Solaris Cluster 4.3 documentation.

After update to Oracle Solaris Cluster 4.3.4, no reconfiguration is necessary for protection groups that were created with the `replication_mode` property in previous versions. The property is ignored by Geographic Edition commands and is no longer listed in configuration output. This change has no effect on the protection group's behavior.

- Oracle Solaris Cluster 4.3.4 software introduces version 6 of the `SUNW.LogicalHostname` resource type. The following changes apply to resources that are created with, or upgraded to, version 6 of the `SUNW.LogicalHostname` resource type:
 - The `Global_zone` property is set to `False` by default.
 - The `Global_zone_override` property is no longer supported. RGM methods will be executed in the zone where the resource group is configured.
- Beginning with Oracle Solaris Cluster 4.3.4 software, the `Active_data_guard` extension property is added to the `SUNW.oracle_server` resource type. The extension property's characteristics are as follows:

`Active_data_guard` (Boolean)

Indicates whether a physical standby instance is additionally managed by Oracle Active Data Guard. Set this property to `TRUE` if the database instance being managed is in an Oracle Active Data Guard configuration. If a `SUNW.oracle_server` resource that is configured with the `Active_data_guard` extension property is restarted on the standby cluster, the Oracle database is started in read-only mode.

Default: `False`

Range: None

Tunable: When Disabled

- The `-G global-devices-file-system` option to the `clnode` command is no longer used and has been removed from the `clnode` command. When adding a node to the cluster, a file

system will automatically be created from a `lofi` device for the global device name space. Removing a node from the cluster removes the file system created from the `lofi` device on that node.

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 StorageTek QFS

At the time of initial release, Oracle Solaris Cluster 4.3 software does not support StorageTek QFS software. Contact your Oracle support representative to learn whether a version of StorageTek QFS software becomes supported with Oracle Solaris Cluster 4.3. You can also check the [Oracle Solaris Cluster 4 Compatibility Guide](#) for the latest StorageTek QFS support information.

Restriction of the Use of EMC SRDF in a Campus Cluster

At the time of initial release, the use of EMC Symmetrix Remote Data Facility (SRDF) data replication for a campus cluster configuration is not yet qualified with Oracle Solaris Cluster 4.3 software. The Oracle Solaris Cluster 4.3 documentation that references this configuration should currently be ignored. Consult the [Oracle Solaris Cluster 4 Compatibility Guide](#) to learn whether this capability later becomes qualified with Oracle Solaris Cluster 4.3 software.

FIPS-140 Compatibility

Most data services are capable of being used on a cluster node running Oracle Solaris 11.3 in FIPS-140 mode. However, the following data services are not compatible with FIPS-140:

- HA for Oracle Business Intelligence Enterprise Edition
- HA for Oracle Database
- HA for Oracle E-Business Suite
- HA for Oracle GlassFish Server Message Queue
- HA for Oracle iPlanet Web Server
- HA for Samba
- HA for Siebel
- HA for Sybase ASE

For the latest information about which data services are not FIPS-140 compatible, see the [Oracle Solaris Cluster 4 Compatibility Guide](#).

The Output of the `verifyrwp` projects API Was Changed From the Original Format (23108669)

Problem Summary: The output of the `verifyrwp` projects API in the ZFS Storage Appliance AK releases are undeterministic and might produce outputs in different formats. Because of this, the `clnasdevice add-dir` command might add ZFS Storage Appliance projects that do not exist on the Storage Appliance device or projects that exist but have invalid configuration to the cluster configuration, but the command would report no error messages about such invalid projects.

Workaround: Before you use the `clnasdevice add-dir` command to add any projects to the cluster configuration, run the following command to identify valid ZFS Storage Appliance projects:

```
# clnasdevice find-dir zfssa-device-name
```

Restriction of Oracle Grid Infrastructure Support for DLMP Link Aggregations or VNICs in Shared-IP Zones (21660315)

Oracle Grid Infrastructure software does not currently support the use of DLMP link aggregations or VNICs in a shared-IP non-global zone. This restriction affects Oracle Solaris Cluster zone cluster configurations with Oracle RAC using Grid Infrastructure, and with HA for Oracle Database using Oracle ASM.

To use Grid Infrastructure in an Oracle Solaris Cluster configuration running Oracle RAC or HA for Oracle Database in a shared-IP zone cluster, use only IPMP groups for public network management.

md_stripe: WARNING: md: write error, md: Panic due to lack of DiskSuite state (21785654)

Problem Summary: Sometimes newfs can cause panic in iSCSI LUNs with ZFS Storage Appliance (ZFSSA) 2011.1.9.x firmware and SVM metaset configuration.

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

listrwprojects Client Interface Not Listing Projects That Have Extra IPs (19982694)

Problem Summary: The `clnas find-dir` command fails to display ZFSSA projects that are setup for use by the cluster when the project contains additional IPs in its NFS Exceptions list. This problem is not encountered if the project's NFS Exceptions list only contains the IPs corresponding to the cluster nodes.

Workaround: Remove the extra IPs in the NFS Exceptions list if the extra IPs are not required. If you require the extra IPs in the NFS Exception list, add the project using the `clnas add-dir project` command.

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.

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 zoneadm 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 4.3. Contact Oracle support services to find out whether a code fix has become available.

Installation Issues

Same DID Device Being Created for All Boot Disks When Installed With Oracle VM Server for SPARC Template (23755653)

Problem Summary: After deploying the Oracle VM Server for SPARC template for Oracle Solaris Cluster when using a boot disk that is backed by a file or a ZFS volume (`zvol`), the DID number for the boot disk is the same on all nodes. Such devices might then be mistakenly considered available for certain operations, like configuring a quorum device.

Workaround: Configure a raw-disk device as the boot disk on all target domains being deployed. For example:

```
# /opt/ovmtutils/bin/ovmtdeploy -d osc43 \  
-o /domains/osc43 -k -s -c 8 \  
-e net0,net2,net3 \  
-v /dev/rdisk/c0t5000CCA00AC0E10Cd0s2 \  
osc-template-location
```

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`

- nlge
- nlsp
- 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)

Administration Notes

This chapter contains known issues and bugs that affect the administration of Oracle Solaris Cluster 4.3. Contact Oracle support services to find out whether a code fix has become available.

Administration Issues

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.

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.

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

Problem Summary: In a campus cluster, 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.

clzonecluster apply Fails to Add Device and Filesystem Into the shared-ip Zone with Errors (21541048)

Problem Summary: When IPv6 or ACFS is configured in a zone, the `clzonecluster apply` command might fail with an error and without making any configuration changes to the system. This issue occurs because the `zonecfg -z zc info -r` command fails to list live configurations of the zone.

Workaround: When IPv6 or ACFS is configured in a zone, reboot the zone cluster instead of using the `clzonecluster apply` command to reconfigure the zone cluster.

Runtime Notes

This chapter contains known issues and bugs that affect the runtime of Oracle Solaris Cluster 4.3. 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.

Oracle Solaris Cluster Manager Browser Interface Notes

This chapter contains known issues and bugs that affect the browser interface of Oracle Solaris Cluster Oracle Solaris Cluster 4.3. Contact Oracle support services to find out whether a code fix has become available.

Browser Interface Issues

HA-Zones Wizard Should Permit Kernel Zones Live Migration for SPARC (23025005)

Problem Summary: When running the HA for Oracle Solaris Zones configuration wizard on a SPARC cluster, the wizard does not permit setting live migration for kernel zones.

Workaround: Use the `clsetup` utility instead. Consult Oracle support to learn if a fix becomes available.

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 browser interface.

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

Oracle Solaris Cluster Manager Configuration Wizard Configures Only Three Nodes Out of the Four (21490228)

Problem Summary: While using the Oracle Solaris Cluster Manager configuration wizard to configure a four-node cluster from one of the nodes, the user interface reports an exception after successfully configuring three nodes. Because the Finish button is disabled, the user cannot configure the last cluster node on the local node.

Workaround 1: Close the browser interface configuration wizard and use `scinstall` to configure the last node of the cluster.

Workaround 2: When the exception is thrown in the browser interface wizard, close the popup showing the error. The Finish button will be disabled so instead click on the Cancel button. The wizard will proceed to the last panel where the cluster check logs will be displayed. Press the Finish button in this panel and that will start the configuration on the last node.

Oracle Solaris Cluster Manager Browser Interface Cannot Run Under Trusted Extensions (21323252)

Problem Summary: The Oracle Solaris Cluster Manager browser interface cannot be used if Trusted Extensions is enabled.

Workaround: Use the command-line interface for managing clusters with Trusted Extensions.

Oracle Solaris Cluster Manager Browser Interface Reports a Connection to the Server Failed Error (21480830)

Problem Summary: A pop-up window with the error, connection to the server failed, shows up. The error status is 500.

Workaround: Close the existing browser window and login again to the user interface.

Deleting Zone Cluster by Using `clzonecluster` Causes the Oracle Solaris Cluster Manager Page to Loop (21555137)

Problem Summary: If the zone cluster details page is open in the Oracle Solaris Cluster Manager browser interface when the same zone cluster is being deleted by using the `clzonecluster delete` command, then the browser interface keeps refreshing indefinitely.

Workaround: Close all tabs of the browser to end the session. If required, manually delete the session cookies. Restart the browser interface.

Errors DCA-29000, http 500, JBO-29114 Displayed If User Removes the Current Cluster From Site on Site Protection Group (21661908)

Problem Summary: An error message is displayed on the Site Details page if you remove the current cluster from the site.

Workaround: Perform the following workaround steps:

1. Remove the cookie for the node hosting the application server from the browser's saved cookies.
2. Log in again to the same node through the browser.
3. Remove the cluster from the site in one of the following ways:
 - Remove the current cluster from the site from the All Sites table by using the Leave Site action, rather than the Site Details page.
 - Remove the current cluster from the site by issuing commands at the command line.

Data Services Notes

This chapter contains known issues and bugs that affect data services of Oracle Solaris Cluster Oracle Solaris Cluster 4.3. 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).

HASP Resources Fail in Oracle Solaris Cluster 4.3 on Oracle Solaris 11.2 and Oracle Solaris 11.3 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.

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, the HASP resources might fail in Oracle Solaris Cluster 4.3 running on Oracle Solaris 11.2 or Oracle Solaris 11.3.

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 HASP resource:

```
# clresource disable hasp-resource-name
```

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

```
# clresource unmonitor hasp-resource-name
```

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

Note that even with the workaround, if a failover of HASP happens during `zfs receive`, the snapshot replication will not complete. You must manually resume the replication on the node that the HASP fails over to.

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.

NFS Server Failover Triggers Stale NFS File Handle (21459179)

Problem Summary: When you reboot or shutdown a cluster node where the Oracle Solaris Cluster HA for NFS resource is online, if an NFS client had an open file or directory which is under write operation, the NFS client might see the Stale NFS file handle error.

Workaround: Before you reboot or shut down the cluster node where the Oracle Solaris Cluster HA for NFS resource is online, execute a resource group switchover to a different target cluster node.

```
# clrg switch -n target_host nfs-rg
```

where `target_host` is the target cluster node for the switchover of resource group `nfs-rg`.

Upgrading From Oracle Solaris 11.2 to Oracle Solaris 11.3 Results in Oracle Grid 12.1.0.1.0 Startup Hang (21511528)

Problem Summary: Oracle Grid startup might hang indefinitely when using Oracle Solaris 11.3 and Oracle Grid 12.1.0.1.0.

Workaround: You can use Oracle Grid 12.1.0.2.0 or Oracle Solaris 11.2 to avoid this problem. Contact Oracle support representative to learn whether a workaround or fix is available.

ORA-00742: Log Read Detects Lost Write (21186724)

Problem Summary: When using Oracle Solaris Cluster HA for Oracle with Solaris Volume Manger (SVM) or UFS file system devices in an x64 cluster environment, Oracle Database log corruption might occur.

Workaround: To avoid data corruption when using SVM or UFS based file systems with HA for Oracle database, place the Oracle binaries and Oracle data on separate file systems. In Oracle data file systems, set `forcedirectio` in `/etc/vfstab` to avoid the bug. You must use `forcedirectio` only for the Oracle data file system, thus requiring separate file systems for Oracle binaries and Oracle data.

Developer Environment Notes

This chapter contains known issues and bugs that affect the developer environment of Oracle Solaris Cluster Oracle Solaris Cluster 4.3. 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 (18528191)

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 4.3 software. Contact Oracle support services to find out whether a code fix has become available.

Geographic Edition Issues

Case Where DIDs Get Reassigned in a Geographic Edition Oracle ZFS Storage Appliance Configuration (24851015)

Problem Summary: Because Oracle ZFS Storage Appliance replicated LUNs cannot be exported while a cluster is secondary for an Oracle ZFS Storage Appliance protection group, the cluster cannot access these LUNs. If the user runs the `cldevice clear` command on this secondary cluster, the DIDs corresponding to these replicated LUNs are removed. If the user then adds new LUNs to the secondary cluster and runs the `cldevice populate` command, DIDs that had been assigned to the deleted replicated LUNs might get reassigned to newly added LUNs.

If later the cluster becomes primary, when the application that uses this replicated data starts on this cluster, attempts to access DIDs that had been assigned to deleted replicated LUNs will not find the expected data, and the application will fail to start.

Workaround: To avoid this issue, never run the `cldevice clear` command on a cluster that is *secondary* for an Oracle ZFS Storage Appliance protection group.

If you encounter this problem, you can use the `cldevice rename` command to resolve the issue. The following scenario illustrates one instance of this problem and the commands to recover from it. The scenario uses the following example components:

- `clusterB` – The secondary cluster for the Oracle ZFS Storage Appliance protection group.

- `zfssaBoxB` – The current target for the replicated project.
- DID 13 – The DID on `clusterB` that corresponds to the replicated LUN in the project that is managed by this protection group.

The following series of actions would create this problem:

1. Add a new LUN to `clusterB`.
2. On one node of `clusterB`, issue the `cldevice clear` command.
DID 13 is removed from the cluster, since the replicated LUN is not exported and cannot be accessed.
3. On one node of `clusterB`, issue the `cldevice populate` command.
DID 13 is assigned to the new LUN created in Step 1.
4. Switch over the protection group to `clusterB` to make `clusterB` the primary.
The switchover issues the `cldevice populate` command. The cluster allocates the next available DID, 14, to the replicated LUN that is now accessible on `clusterB`.

The application resource is now unable to start, because the data in DID 13 is not what is expected.

The following recovery steps correct the problem, where DID 15 is the next unassigned DID:

1. On each node of `clusterB`, move DID 13 to DID 15.

```
# cldevice rename -d 15 13; devfsadm; cldevice populate
```
2. On each node of `clusterB`, move DID 14 to DID 13.

```
# cldevice rename -d 13 14; devfsadm; cldevice populate
```
3. Restart the application resource group.

```
clresourcegroup restart rg-name
```

The application can now start because it finds the expected data in DID 13.

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
```

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.

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.

Infrastructure Resource Goes Offline After Probe Caused a Restart and Failed to Start (21298474)

Problem Summary: If the Geographic Edition setup on a cluster has multiple protection groups and multi-group configurations, the related infrastructure components might take a long time to start. This startup is managed by the `geo-failovercontrol` resource of the `SUNW.scmasa` resource type, which has a default start timeout of 600 seconds. If the `geo-failovercontrol` resource takes more time to start than the default start timeout, the Geographic Edition infrastructure goes offline.

Workaround: Increase the `Start_timeout` property value of the `geo-failovercontrol` resource in the `geo-infrastructure` resource group. If the `RG_system` property of the `geo-infrastructure` resource group is `TRUE`, temporarily change it to `FALSE` before changing the resource property.

Type the following commands to change the `Start_timeout` of the resource to 1200 seconds.

```
$ /usr/cluster/bin/clresourcegroup set -p RG_system=FALSE geo-infrastructure
$ /usr/cluster/bin/clresource set -p Start_timeout=1200 geo-failovercontrol
$ /usr/cluster/bin/clresourcegroup set -p RG_system=TRUE geo-infrastructure
```

Oracle GoldenGate Protection Group Data Replication Status Shows OK Though the Replication Resource did not Failover (21527062)

Problem Summary: After a node failure on the secondary partner, the Oracle GoldenGate replication status resource does not start up on another node of the secondary partner because the affinity resource group of the Oracle GoldenGate replication status resource group did not come up. This behavior is valid as per the resource group affinity. However, the protection group's data replication status does not reflect the replication status resource's new status and the replication status still shows OK.

Workaround: Validate the protection group using `geopg validate` on the cluster, which would query the latest replication resource status and update the protection group's replication status.

java.lang.IllegalArgumentException: Unmatched braces in the pattern (21570583)

Problem Summary: Protection group creation fails if one of the cluster nodes is down or the common agent container is not running on a node and displays the following error message on the terminal:

```
Cannot reach management agent on cluster-node :  
Internal Error :javax.management.RuntimeMBeanException:  
java.lang.IllegalArgumentException: Unmatched braces in the pattern.
```

Workaround: Ensure that the common agent container is running on all cluster nodes. If a node is down, bring up the node or remove the node and create the protection group.

Protection Group Creation Should Not Fail If a Node in the Cluster Is Down (21697993)

Problem Summary: Protection group creation fails if one of the cluster nodes is down. This situation occurs when the script-based plug-in module tries to check whether all *_script files exist and are executable on all cluster nodes. The check is performed on all nodes because the script-based plug-in module does not have a script-based plug-in name to look up in the configuration file. If one of the cluster nodes is down, then an exception is thrown which terminates the protection group creation.

Workaround: Bring up the node or remove the node and create the protection group.

If Takeover Is Performed While Both Sites Are Up, Project Is Not Removed From the Original Primary Site (21684495)

Problem Summary: If you run the `geopg takeover` command when both the primary and secondary ZFSSA appliances are up, then switchover to the secondary site fails because of an empty project that exists on the original primary ZFSSA appliance after the protection group is activated.

Workaround: Before attempting to switchover the protection group, remove the empty project on the secondary appliance after the protection group is activated.

Geographic Edition Does Not Support ZFSSA Offline Replication Feature (21934145)

Problem Summary: Geographic Edition incorrectly allows a switchover while replication is in the Idle (export pending) state.

Workaround: Do not use offline replication feature on projects managed by the Geographic Edition.

Documentation Notes

This chapter discusses errors or omissions for documentation in the Oracle Solaris Cluster 4.3 release.

Man Pages

This section provides the following information about errors, omissions, and additions in the Oracle Solaris Cluster man pages:

- [“Extraneous Characters in the Command Lines of Some Man Page Examples” on page 51](#)
- [“geopg\(1M\) Man Page” on page 51](#)

Extraneous Characters in the Command Lines of Some Man Page Examples

Some command examples contain extraneous letters after the backslash (\) used to indicate that the command continues on the next line. For example:

```
phys-schost-1# /usr/cluster/bin/clresource create epm-rg \fR
-t SUNW.HAStoragePlus \fR
-p FileSystemMountPoints=/global/epm_mnt \fR
-d epm-hasp-r
```

The fR letters after the backslash can be ignored.

geopg(1M) Man Page

- Beginning with the Oracle Solaris Cluster 4.3.6 release, the `geopg status` command supports the following additional option:

-v Displays messages for any errors that put the configuration in the Error state, and messages for any errors that put the replication resource in the FAULTED or DEGRADED error states.

--verbose

- The geogg (1M) man page is missing information about the extension properties for Oracle GoldenGate data replication for Geographic Edition. See [Appendix A, “Oracle GoldenGate Replication Extension Properties,”](#) in *Oracle Solaris Cluster Geographic Edition Data Replication Guide for Oracle GoldenGate* for information about these extension properties.
- Descriptions of some geogg subcommands erroneously state that the action it performs is propagated to clusters in a site, similar to the following:

The system performs this action on the local cluster, then propagates the action to other clusters in the site.

The statement should instead say that the action is propagated to the partner cluster. For example:

The system performs this action on the local cluster, then propagates the action to the partner cluster.

New and Enhanced Extension Properties for the SUNW.Proxy_SMF_* Resource Types

Beginning with Oracle Solaris Cluster 4.3.4 software, the following extension properties are added to the SUNW.Proxy_SMF_failover, SUNW.Proxy_SMF_multimaster, and SUNW.Proxy_SMF_scalable resource types.

Monitor_retry_count Number of PMF restarts allowed for the fault monitor.

Default: 2

Tunable: Any time

Monitor_retry_interval Time window (in minutes) for fault monitor restarts.

Default: 2

Tunable: Any time

Probe_command Name of the command to be run by the SMF proxy probe.

Default: " "

Tunable: When disabled

Probe_timeout Timeout value for the probe (in seconds).

Default: 30

Minimum: 2

Tunable: Any time

In addition, the `Proxied_service_instances` extension property can now be specified directly to the `clresource` command with the `-x` option. The following is the revised description of this property.

`Proxied_service_instances` Specifies information about the SMF services to be proxied by the resource. Its value is either one or more `fmri` names for the SMF services to proxy or the path to a file that contains all the proxied SMF services.

- The preferred method to specify this property is by direct use in the `clresource` command. If using the `fmri` name, use the following syntax:

```
-x Proxied_service_instances=svc:service-name
```

You can specify multiple SMF services in a comma-delimited list. For example:

```
-x Proxied_service_instances=svc:/network/ntp:default,svc:/network/dhcp-server:default
```

- The use of a file to provide the `Proxied_service_instances` values is deprecated but still supported. When specifying the property in a file, the property's value is the path to a file that contains all the proxied SMF services. Each line in the file is dedicated to one SMF service and specifies `svc fmri` and the path to the corresponding service manifest file. For example, if the resource has to manage two services, `restarter_svc_test_1:default` and `restarter_svc_test_2:default`, the file must include the following two lines:

```
<svc:/system/cluster/restarter_svc_test_1:default>,  
</var/svc/manifest/system/cluster/restarter_svc_test_1.xml>
```

```
<svc:/system/cluster/restarter_svc_test_2:default>,  
</var/svc/manifest/system/cluster/restarter_svc_test_2.xml>
```

Note - The entries above must each appear on a single line. They are broken into multiple lines here for legibility.

Default: ""

Tunable: When disabled

Data Services Planning and Administration Guide

- The procedure [“How to Configure a Failover Application Using the ScalMountPoint Resource”](#) in *Oracle Solaris Cluster 4.3 Data Services Planning and Administration Guide* is missing the following step, to optionally perform after Step 1:

2. If the application binaries use a NAS NFS file system and you want the resource to automatically fail over if a storage fault is detected, set the `RebootOnFailure` property to `True`.

Setting this property prevents the resource from entering a `STOP_FAILED` state if storage connectivity is faulted. Instead, the node where the `ScalMountPoint` resource resides is rebooted, and the resource restarts on another cluster node.

Note - The availability of other services might be adversely impacted by this fault case. Before you set this property, ensure that you have considered its impact to all services running on the cluster node. You can limit the impact of the `RebootOnFailure` setting on other services by instead configuring this service in a zone cluster. The reboot then only affects services within that zone cluster.

```
# clresource set -p RebootOnFailure=True scalable-mount-point-resource
```

- The procedure [“How to Configure a Scalable Application Using the ScalMountPoint Resource”](#) in *Oracle Solaris Cluster 4.3 Data Services Planning and Administration Guide* is missing the following step, to optionally perform after Step 1:

2. If the application binaries use a NAS NFS file system and you want the node where the resource was online to abort if a storage fault is detected, set the `RebootOnFailure` property to `True`.

Setting this property prevents the resource from entering a `STOP_FAILED` state if storage connectivity is faulted. This `STOP_FAILED` state might leave application instances in a hanging state on the node that has the fault and possibly also on other nodes if the instances communicate among themselves. But if the `RebootOnFailure` property is set to `True`, the node where the `ScalMountPoint` resource resides would instead abort. Manual intervention is then necessary to repair the storage connection for the faulted node and restart the storage resource on that node, which then restarts the application instances on that node.

Note - The availability of other services might be adversely impacted by this fault case. Before you set this property, ensure that you have considered its impact to all services running on the cluster node. You can limit the impact of the `RebootOnFailure` setting on other services by instead configuring this service in a zone cluster. The reboot then only affects services within that zone cluster.

```
# clresource set -p RebootOnFailure=True scalable-mount-point-resource
```

HA for Oracle External Proxy Guide

When creating a new resource for the HA for Oracle External Proxy data service, do not follow instructions in the data service guide to use the `openssl` command to encrypt the password for the remote Oracle database user. For Oracle Solaris Cluster 4.3, this password encryption is now performed by using the Oracle Solaris Cluster private string command, `clpstring`. Issue the following command from one node of the cluster:

```
# clpstring create -b resource resource-pw
Enter string value: *****
Enter string value again: *****
#
```

Note that the naming convention for the private string is the name of the resource appended with "-pw". For example, for a resource named `oep-proxy-rs`, you name the private string `oep-proxy-rs-pw`.

The 4.3 version of the agent automatically converts any existing resource that uses the `openssl` encrypted password to use a private string instead. The agent also removes the `openssl` password files, but not the `openssl` key files. You can remove these key files once you no longer use them. For more information about using FIPS 140 cryptography, see [Using a FIPS 140-2 Enabled System in Oracle Solaris 11.3](#).

Note - If, before the agent automatically converts the password, you converted the cluster node to use the FIPS-140-capable `openssl` command, the agent will fail to decrypt the old password and will not be able perform the automatic conversion. In this situation, use the above `clpstring` command to create the private string.

HA for PostgreSQL Guide

When you encrypt a password do not run the encryption on all nodes. Run the encryption on one node only. This requirement must be observed in the following procedures:

- [“How to Create and Enable Resources for PostgreSQL” in Oracle Solaris Cluster Data Service for PostgreSQL Guide](#)
- [“Example: Enabling the PostgreSQL Software to Run in the Cluster” in Oracle Solaris Cluster Data Service for PostgreSQL Guide](#)
- [“Example: Enabling the PostgreSQL Software to Run in the Cluster” in Oracle Solaris Cluster Data Service for PostgreSQL Guide](#)

Hardware Administration Manual

Corrections and clarifications about supported data replication methods in a campus cluster were made to [“Requirements for Designing a Campus Cluster” in Oracle Solaris Cluster Hardware Administration Manual](#). These changes were published as of February 2017.

Software Installation Guide

Oracle Solaris Cluster documentation omits a restriction against configuring a SUNW.SharedAddress resource in an exclusive-IP zone cluster. To configure a SUNW.SharedAddress resource in a zone cluster, the zone cluster must be of the shared-IP type.

System Administration Guide

Corrections and clarifications about support for data replication products in a campus cluster or Geographic Edition configuration were made to [Chapter 4, “Data Replication Approaches” in Oracle Solaris Cluster 4.3 System Administration Guide](#). These changes were published as of February 2017.

Addendum to Support Network Appliance NAS Devices

The following supplementary information supports the use of the `clnasdevice` command to configure a Network Appliance NAS device:

-d directory

Use to specify the directory or directories of the netapp NAS device.

-t nas-device-type

netapp

Specifies a NAS device from Network Appliance, Inc. The NAS device from Network Appliance, Inc. has the following property. This property is required if you add a NAS device by using the add subcommand:

`-p userid=userid [-f passwdfile]`

or

`-u userid [-f passwdfile]`

The *userid* is the user ID that the cluster uses to perform administrative duties on the NAS device. When you add a user ID to the device configuration, you are prompted for its password. You can also place the password in a text file and use it by specifying the `-f` option.

To provide support for netapp NAS devices in the cluster, the administrator must also install the required software module that is provided by Network Appliance, Inc. Additionally, the iSCSI license must be valid for the Network Appliance, Inc. NAS device. For instructions about obtaining the support module, see [Oracle Solaris Cluster With Network-Attached Storage Device Manual](#).

