

Oracle Solaris Cluster 4.4 Release Notes



E69309-06
February 2024



Oracle Solaris Cluster 4.4 Release Notes,

E69309-06

Copyright © 2000, 2024, Oracle and/or its affiliates.

Primary Author: Cathleen Reiher, Kristi Caruso

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, software documentation, data (as defined in the Federal Acquisition Regulation), 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," "commercial computer software documentation," or "limited rights data" 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®, Java, MySQL and NetSuite 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.

Copyright © 2000, 2024, Oracle et/ou ses affiliés.

Ce logiciel et la documentation connexe sont fournis en vertu d'un contrat de licence assorti de restrictions relatives à leur utilisation et divulgation. Ils sont protégés en vertu des lois sur la propriété intellectuelle. Sauf dispositions contraires prévues de manière expresse dans votre contrat de licence ou permises par la loi, vous ne pouvez pas utiliser, copier, reproduire, traduire, diffuser, modifier, mettre sous licence, transmettre, distribuer, présenter, effectuer, publier ou afficher à toutes fins une partie de ces derniers sous quelque forme que ce soit, par quelque moyen que ce soit. Sont interdits l'ingénierie inverse, le désassemblage ou la décompilation de ce logiciel, sauf à des fins d'interopérabilité selon les dispositions prévues par la loi.

L'information contenue dans les présentes est sujette à changement sans préavis. Nous ne garantissons pas qu'elle est exempte d'erreur. Si vous y relevez des erreurs, veuillez nous les signaler par écrit.

Si ce logiciel, la documentation du logiciel ou les données (comme défini dans la réglementation Federal Acquisition Regulation) ou la documentation afférente sont livrés sous licence au gouvernement des États-Unis d'Amérique ou à quiconque qui aurait souscrit la licence de ce logiciel pour le compte du gouvernement des États-Unis d'Amérique, la notice suivante s'applique :

UTILISATEURS DE FIN DU GOUVERNEMENT É.-U. : programmes Oracle (y compris tout système d'exploitation, logiciel intégré, tout programme intégré, installé ou activé sur le matériel livré et les modifications de tels programmes) et documentation sur l'ordinateur d'Oracle ou autres logiciels Oracle. Les données fournies aux utilisateurs finaux du gouvernement des États-Unis ou auxquelles ils ont accès sont des "logiciels informatiques commerciaux", des "documents sur les logiciels informatiques commerciaux" ou des "données relatives aux droits limités" conformément au règlement fédéral sur l'acquisition applicable et aux règlements supplémentaires propres à l'organisme. À ce titre, l'utilisation, la reproduction, la duplication, la publication, l'affichage, la divulgation, la modification, la préparation des œuvres dérivées et/ou l'adaptation des i) programmes Oracle (y compris tout système d'exploitation, logiciel intégré, tout programme intégré, installé, ou activé sur le matériel livré et les modifications de ces programmes), ii) la documentation informatique d'Oracle et/ou iii) d'autres données d'Oracle, sont assujetties aux droits et aux limitations spécifiés dans la licence contenue dans le contrat applicable. Les conditions régissant l'utilisation par le gouvernement des États-Unis des services en nuage d'Oracle sont définies par le contrat applicable à ces services. Aucun autre droit n'est accordé au gouvernement américain.

Ce logiciel ou matériel informatique est destiné à un usage général, dans diverses applications de gestion de l'information. Il n'a pas été conçu pour être utilisé dans le cadre d'applications dangereuses, y compris des applications susceptibles de causer des blessures corporelles. Si vous utilisez ce logiciel ou matériel informatique dans des applications dangereuses, il vous revient d'adopter les mesures relatives à la protection contre les interruptions, aux copies de sauvegarde et à la redondance ainsi que toute autre mesure visant à garantir son utilisation en toute sécurité. Oracle Corporation et ses sociétés affiliées déclinent toute responsabilité relativement aux dommages pouvant résulter de l'utilisation du logiciel ou du matériel informatique dans des applications dangereuses.

Oracle®, Java, MySQL et NetSuite sont des marques de commerce enregistrées d'Oracle Corporation et/ou de ses sociétés affiliées. Les autres noms ou raisons sociales peuvent être des marques de commerce de leurs propriétaires respectifs.

Intel et Intel Inside sont des marques de commerce ou des marques de commerce enregistrées de Intel Corporation. Toutes les marques de commerce SPARC sont utilisées sous licence et sont des marques de commerce ou des marques de commerce enregistrées de SPARC International, Inc. AMD, Epyc et le logo AMD sont des marques de commerce ou des marques de commerce enregistrées de Advanced Micro Devices. UNIX est une marque de commerce enregistrée de The Open Group.

Ce logiciel ou matériel informatique et sa documentation peuvent fournir de l'information sur du contenu, des produits et des services tiers, ou y donner accès. Oracle Corporation et ses sociétés affiliées déclinent toute responsabilité quant aux garanties de quelque nature que ce soit relatives au contenu, aux produits et aux services offerts par des tiers, sauf mention contraire stipulée dans un contrat entre vous et Oracle. Oracle Corporation et ses sociétés affiliées ne pourront être tenus responsables des pertes, frais et dommages de quelque nature que ce soit découlant de l'accès à du contenu, des produits ou des services tiers, ou de leur utilisation, sauf mention contraire stipulée dans un contrat entre vous et Oracle.

Contents

Using This Documentation

Product Documentation Library	viii
Feedback	viii
Diversity and Inclusion	viii

1 Oracle Solaris Cluster 4.4 What's New

What's New in the Software	1-1
Former Features Not Included in the Oracle Solaris Cluster 4.4 Software	1-2
Man Page Section Renumbering	1-3

2 Compatibility Notes

Compatibility Issues	2-1
Restriction of Oracle Grid Infrastructure Support for DLMP Link Aggregations or VNICs in Shared-IP Zones (21660315)	2-1
Zone Does Not Boot if pkg:/system/resource-mgmt/resource-caps Is Not Installed and capped-memory Is Configured (15740089)	2-1
Upgrading From Oracle Solaris 11.2 to Oracle Solaris 11.4 Results in Oracle Grid 12.1.0.1.0 Startup Hang (21511528)	2-1
SUNW.qfs Mount Fails if File Name Matching special-device Name Exists Under "/" (30333778)	2-2
Oracle Solaris Cluster Manager Login to Cluster Node Fails Starting With Oracle Solaris 11.4 SRU 33 (32871272)	2-2
Solaris Firewall Service Prevents Oracle Solaris Cluster 4.4 From Rebooting in to Cluster Mode (32987543)	2-3

3 Installation Notes

Installation Issues	3-1
rpc-authdes Utilities Moved From core-os To legacy/security/rpc-authdes (30680102)	3-1
Hard to Determine Data Service Names for solaris10 Branded Zone Noninteractive Data Service Installation (15804349)	3-1

4 Administration Notes

Administration Issues	4-1
Depending on vfstab Creation of nfs rs Fails w INTERNAL ERROR if Failover zpool (28546356)	4-1
cacaoadm start Fails With Warning: found remaining processes in SMF contract (28508074)	4-1
Restrictions for SVMs in Zone Cluster (28384238, 28313846)	4-2
Several Cluster Checks Fail With Execution Error (27092918)	4-2
zoneadm Commands Hanging at Bootup into 11.4, Cluster 4.4 Following Upgrade (25743399)	4-2
clzonecluster apply Fails to Add Device and Filesystem Into the shared-ip Zone with Errors (21541048)	4-2
Failed to Remove a File from Global ZFS Filesystem When Quota is Exceeded (20860157)	4-2
clzonecluster install -a archive-no-cluster-pkgs zone-cluster Does Not Install Cluster Packages (18714803)	4-3
Using chmod to Set setuid Permission Returns Error in a Non-Global Zone on PxFs Secondary Server (15697690)	4-3
IPv6 Scalable Service Support Is Not Enabled By Default (15290321)	4-3
Removing a Node From an Exclusive-IP Zone Cluster Panics the Cluster Nodes (15817184)	4-4

5 Runtime Notes

Runtime Issues	5-1
Investigate Write Performance Improvements for Global ZFS (26201389)	5-1

6 Oracle Solaris Cluster Manager Browser Interface Notes

Browser Interface Issues	6-1
Selecting or Unselecting Accessibility Item Results in Blank Screen (28003165)	6-1
On Newly Booted Nodes Sometimes Browser Becomes Unresponsive (26675695)	6-1
Tasks - Highly Available Storage Can Not Proceed in ES (23135263)	6-1
Oracle Solaris Cluster Manager Configuration Wizard Configures Only Three Nodes Out of the Four (21490228)	6-1
Oracle Solaris Cluster Manager Browser Interface Cannot Run Under Trusted Extensions (21323252)	6-2
Unable to Install Oracle Solaris Cluster for solaris10 Branded Zone (19064831)	6-2
RAC Wizard Fails w java.lang.NullPointerException When Creating Clusterware rs (27553809)	6-2
Oracle DB GUI Wizard Does Not Show Any ASM Projects Leading to Failure (27548953)	6-2

ADF Error Listing RG's For Add to PG Table (25065856)	6-2
---	-----

7 Data Services Notes

Data Services Issues	7-1
Resource Creation Using SUNW.scalable_rac_server_proxy fails - Specified DB_NAME db_name Not Registered with CRS (29054450)	7-1
Migration of Resources Registered as ORCL.ha-zone_sczbt, ORCL.ha-zone_sczsh or ORCL.ha-zone_sczsmf to the Latest Resource Type Version (21926061)	7-1
HA-LDOMs Resources Must Update to 4.3 before Updating to 4.4 (23308155)	7-2
run_setuid_prog Failing With ORA-01017: Invalid username/password; Logon Denied (21509291)	7-3
Oracle Database/WLS Resource Fails to Come Online Due to Locking Issue (15713853)	7-3
ORA-00742: Log Read Detects Lost Write (21186724)	7-3

8 Disaster Recovery Framework Notes

Disaster Recovery Framework Issues	8-1
Update How to Perform a Failback-Switchover on a System That Uses Oracle ZFS Storage Appliance Replication Procedure (36175231)	8-1
ZFS SA Replication Module Does Not Support Cluster File Systems With ZFS (29223365)	8-4
Case Where Node Stays Fenced Off & Later ScalMountPoint rs Fails to Start There (28362456)	8-4
Flag Cases & Disallow Switchovers Accordingly Where Dependent rgs are Not in pg (28004581)	8-5
geopg Returns Java Exception on 4.4 ZC Node While Other Node is Running 4.3.9 (27148373)	8-5
Switchover Fails if Replicated FS Exported & Mountpoint Can be Overwritten (25970289)	8-5
Case Where DIDs Get Reassigned in a Geographic Edition Oracle ZFS Storage Appliance Configuration (24851015)	8-5
Bug 16964103 About Underlying did Not Available Still Exists if pg in ZC (24582738)	8-7
Dataguard Manager User for Container db (c##dgmgr) Fails (24354703)	8-7
Oracle Data Guard Module Incorrectly Flagging SUNW.oracle_server Dependencies in the Single Instance (15818725)	8-7
Collision Problems Should Be Flagged at Protection Group Creation Time (22529981)	8-7
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)	8-8
If Takeover Is Performed While Both Sites Are Up, Project Is Not Removed From the Original Primary Site (21684495)	8-8

Disaster Recovery Framework Does Not Support ZFSSA Offline Replication Feature (21934145)	8-8
---	-----

9 Documentation Notes

clresource create -d Commands Incorrectly Specify Resource Names in Oracle Solaris Cluster Data Service for Oracle E-Business Suite as of Release 12.2 Guide (26401103)	9-1
Refer to Oracle VM Server for SPARC 3.6 Documentation in Oracle Solaris Cluster Data Service for Oracle VM Server for SPARC Guide (33890482)	9-5
clpstring Command in Oracle Solaris Cluster Data Service for Oracle VM Server for SPARC Guide Has Incorrect Parameters (33837717)	9-5
Add parameter_file Property Section to Setting HA for Oracle Database Extension Properties in Oracle Solaris Cluster Data Service for Oracle Database Guide (33131714)	9-6
Update Instructions About Booting All Nodes Into Non-Cluster Mode and Rebooting Nodes in Administering an Oracle Solaris Cluster 4.4 Configuration (31092326)	9-6
Update Database Preparation in Oracle Solaris Cluster Data Service for Oracle GoldenGate Guide (20170637)	9-7
Add How to Add a ZFS Cluster File System to a Zone Cluster (CLI) Task to Adding File Systems to a Zone Cluster in Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4 (30930843)	9-7
Change Instances of StorageTek to HSM in Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4 (30930843)	9-7
Add Clarification About a SAP xserver to How to Register and Configure an SAP xserver Resource in Oracle Solaris Cluster Data Service for SAP MaxDB Guide (21803412)	9-8
Oracle Database on Oracle Solaris Cluster Does Not Require oracle:dba User and Group Settings (32988568)	9-8
SUNW.oracle_server(7) Man Page Missing From Oracle Solaris Cluster 4.4 Reference Manual (31191927)	9-8
Clarify That How to Remove a Node From a Zone Cluster in Administering an Oracle Solaris Cluster 4.4 Configuration Removes the Zone From the Cluster Node (31145133)	9-8
Incorrect Configuration File Name in Oracle Solaris Cluster Data Service for Apache Tomcat Guide (32973775)	9-9
Update Erroneous clpstring Command in Oracle Solaris Cluster Data Service for Oracle VM Server for SPARC Guide (32988568)	9-9
Update Misspelled Property Name in Oracle Solaris Cluster Data Service for Oracle VM Server for SPARC Guide (32988568)	9-9
List of Zone Cluster Tasks (31080689)	9-10
Add Requirements to How to Use Link Aggregation Control Protocol (LACP) with Cluster Interconnects in Administering an Oracle Solaris Cluster 4.4 Configuration (33099699)	9-11
geopg(8) Man Page	9-12

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://docs.oracle.com/cd/E69294_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.

1

Oracle Solaris Cluster 4.4 What's New

This chapter provides the following information about the Oracle Solaris Cluster 4.4 software:

- [What's New in the Software](#)
- [Former Features Not Included in the Oracle Solaris Cluster 4.4 Software](#)
- [Man Page Section Renumbering](#)

Note:

Check the Compatibility Guide for the support of your database versions with Oracle Solaris Cluster 4.4.

What's New in the Software

This section highlights information for existing customers about new features in Oracle Solaris Cluster 4.4 software.

The Oracle Solaris Cluster 4.4 software provides the following new features:

- **Cluster File System with ZFS** – Beginning with Oracle Solaris Cluster 4.4 you can use proxy file system (PxFs) functionality on top of ZFS datasets. This feature facilitates access to mounted ZFS file systems from any node in the Cluster, delivers ZFS functionality for storage management while maintaining POSIX compliance, and extends Cluster Device Configuration System (DCS) to zpool management. See [Creating Cluster File Systems in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*](#).
- **ZFS Snapshot Replication** – Beginning with Oracle Solaris Cluster 4.4 you can use the Oracle Solaris Cluster disaster recovery framework module for Oracle Solaris ZFS snapshot to replicate `zpools` for globally mounted ZFS file systems.

For more information, see [Oracle Solaris Cluster Data Replication Guide for ZFS Snapshots](#).
- **Scalable Services for Exclusive IP Zone Cluster** – Beginning with Oracle Solaris Cluster 4.4, Solaris Cluster Scalable services, which provides load balancing, now works with Exclusive-IP zone cluster (`ip-type=exclusive`). See [Creating and Configuring a Zone Cluster in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*](#).
- **First Boot IP Configuration for Exclusive-IP Zone Clusters** – Beginning with Oracle Solaris Cluster 4.4, exclusive-IP zone clusters can have their public interfaces automatically configured to be on the networks at first boot.
- **Immutable Zone Clusters** – Oracle Solaris Cluster 4.4 now supports deploying HA data services on immutable zone-clusters in fixed-configuration and flexible-configuration profiles. See [Creating and Configuring a Zone Cluster in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*](#).
- **LP64 – Leverage ADI** – Oracle Solaris Cluster 4.4 commands, libraries and plugins are now ported to LP64. Some benefits of a LP64 conversion include:

- Modernization
- Increase performance, especially on x64 platforms where additional CPU registers will be available when the utilities are recompiled in 64-bit mode
- Better randomization with the new Address Space Layout Randomization (ASLR) Oracle Solaris security feature
- An LP64 conversion solves the Year 2038 problem with commands that use or manipulate 32-bit `time_t` values.
- **Oracle WebLogic Server for Oracle Solaris Cluster Manager** – Beginning with Oracle Solaris Cluster 4.4, Oracle GlassFish server is replaced with Oracle WebLogic Server for Oracle Solaris Cluster Manager.
- **Precision Time Protocol (PTP)** – You can now configure Precision Time Protocol (PTP) on Oracle Solaris Cluster 4.4 nodes which otherwise use Network Time Protocol (NTP) by default. If you configure the hardware assistance feature of PTP, the network interface configured for PTP can not be used to host LogicalHostname or ScalableAddress resource IP addresses.
- **Centralized Install** – Beginning with Oracle Solaris Cluster 4.4, you can perform centralized installation of the software. See [How to Install Oracle Solaris Cluster Software Using Centralized Installer \(clinstall\) in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*](#).
- **Disaster Recovery Framework Name**– Beginning with Oracle Solaris Cluster 4.4, to better describe the functionality of the product, the Disaster Recovery feature has been renamed from Geographic Edition to Disaster Recovery Framework.
- **Puppet Support** – Puppet is cross-platform software that you can use to automate and enforce the configuration management of most major subsystems. You can use Puppet to perform several common system configuration tasks. For further details about Puppet refer to [Using Puppet With Oracle Solaris Cluster in *Administering an Oracle® Solaris Cluster 4.4 Configuration*](#) and [Using Puppet to Perform Configuration Management](#).
- **Java 8 Update Considerations** – Updating to Oracle Solaris 11.4 removes Java 7. The default version of Java for Oracle Solaris 11.3 is Java 8. If the default version of Java on your Oracle Solaris 11.3 system was changed by re-setting the java mediated link (as in `pkg set-mediator -V 1.7 java`), then after upgrade to Oracle Solaris 11.4, this link points to a path that does not exist. To use Java 8, do one of the following:
 - Reset the mediated link.


```
# pkg set-mediator -V 1.8 java
```
 - Specify the full path to Java 8: `/usr/jdk/instances/jdk1.8.0/bin/java`
- **Supported Data Services**– For a list of the data services supported by Oracle Solaris Cluster 4.4, see the [Oracle® Solaris Cluster 4 Compatibility Guide](#).

Former Features Not Included in the Oracle Solaris Cluster 4.4 Software

The following features are included in the Oracle Solaris Cluster 4.3 version but are not included in the Oracle Solaris Cluster 4.4 release:

- Support for Veritas header files
- Support for Availability Suite
- Support for RDT
- Support for RSM
- Legacy command set (`sc*`)
- Support for service tags
- Agent Usability in Failover Zones for MySQL, Apache Tomcat, and PostgreSQL

In addition, for features already not included previously in Oracle Solaris Cluster 4.3 see the [Oracle Solaris Cluster 4.3 Release Notes](#).

Man Page Section Renumbering

In Oracle Solaris Cluster 4.4, man pages using the System V sections have been renumbered to the standard sections. The sections 1m, 4, 5, 7, and their subsections used in the previous releases are now 8, 5, 7, 4, and their subsections respectively. Users who are familiar with other platforms such as BSD, Linux, or MacOS X can use the same section numbers with the man command.

In addition, the `libscho.st.so.1(1)` man page has been recategorized as section 3lib.

The man command can map the old System V section numbers to the standard ones when necessary, to help find references from older documentation.

For more information about all man page enhancements, see [man Command Enhancements in What's New in Oracle Solaris 11.4](#) and the [man\(1\)](#) man page.

2

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 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. Or contact an Oracle support representative to learn whether a workaround or fix is available

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

Problem Summary: If the package `pkg:/system/resource-mgmt/resource-caps` 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-caps` package into the global zone. Once the `resource-caps` package is installed, the zone can boot.

Upgrading From Oracle Solaris 11.2 to Oracle Solaris 11.4 Results in Oracle Grid 12.1.0.1.0 Startup Hang (21511528)

Problem Summary: Oracle Grid startup might hang indefinitely when using Oracle Solaris 11.4 and Oracle Grid 12.1.0.1.0.

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

SUNW.qfs Mount Fails if File Name Matching *special-device* Name Exists Under "/" (30333778)

Problem Summary: SUNW.qfs mount fails in Oracle Solaris Cluster 4.4 if the mountpoint for your filesystem is `/<special-device>` or a file or directory exists under `"/"` named `<special-device>` where `<special-device>` is the name of the HSM special device that you are mounting.

Workaround: Change the file or the mountpoint name, if neither the file nor the mountpoint name can be changed, contact an Oracle support representative for a fix.

Oracle Solaris Cluster Manager Login to Cluster Node Fails Starting With Oracle Solaris 11.4 SRU 33 (32871272)

Problem Summary: Starting with Oracle Solaris 11.4 SRU 33, the OS delivers Version `jdk1.8.0.291` of the JDK software. This JDK version causes an attempted login by the Oracle Solaris Cluster manager to a cluster node to fail with the following error:

```
Unable to connect to cluster-node. Ensure that the common agent container is online.
```

The Oracle Solaris Cluster manager and the cluster node to which you want to connect might be on a single machine or on two machines. Either or both of these machines might run at least Version `jdk1.8.0.291` of the JDK software.

Workaround: Perform the following steps based on the machines that use JDK Version `jdk1.8.0.291` in your environment:

- **The Oracle Solaris Cluster manager and the cluster node:** Perform the following steps:
 1. On the machine that runs the Oracle Solaris Cluster manager, remove the `TLSv1` value from the `jdk.tls.disabledAlgorithms` property in the `/usr/jdk/instances/jdk1.8.0/jre/lib/security/java.security` file.

For example, remove the `TLSv1` value from the following property assignment:

```
jdk.tls.disabledAlgorithms=SSLv3, TLSv1, TLSv1.1, RC4, DES, MD5withRSA, \
```

To result in the following property assignment:

```
jdk.tls.disabledAlgorithms=SSLv3, TLSv1.1, RC4, DES, MD5withRSA, \
```

2. On the cluster node to which you want to connect, remove `TLSv1` from the `jdk.tls.disabledAlgorithms` property value in the `/usr/jdk/instances/jdk1.8.0/jre/lib/security/java.security` file.

For example, remove the `TLSv1` value from the following property assignment:

```
jdk.tls.disabledAlgorithms=SSLv3, TLSv1, TLSv1.1, RC4, DES, MD5withRSA, \
```

To result in the following property assignment:

```
jdk.tls.disabledAlgorithms=SSLv3, TLSv1.1, RC4, DES, MD5withRSA, \
```

3. On the machine that runs the Oracle Solaris Cluster manager, restart the `svc:/system/cluster/manager-wls:default` SMF service.

```
# svcadm restart svc:/system/cluster/manager-wls:default
```
 4. On the cluster node to which you want to connect, restart `cacao`.

```
# cacaoadm restart
```
- **Only the Oracle Solaris Cluster manager machine:** Perform Steps 1 and 3.
 - **Only the cluster node:** Perform Steps 2 and 4.

Solaris Firewall Service Prevents Oracle Solaris Cluster 4.4 From Rebooting in to Cluster Mode (32987543)

Description: A cluster boot hangs indefinitely after configuring the Solaris Firewall service for a cluster that runs Oracle Solaris Cluster 4.4.

The following output shows that when you reboot the `node01` node, it never joins the cluster:

```
root@node01:~ May 21 09:42:14 node01 cl_runtime: NOTICE: CMM: Cluster doesn't
have operational quorum yet; waiting for quorum.
May 21 09:42:14 node01 cl_runtime: NOTICE: clcomm: Path node01:net3 -
node02:net1 errors during initiation
May 21 09:42:14 node01 cl_runtime: NOTICE: clcomm: Path node01:net0 -
node02:net0 errors during initiation
May 21 09:42:14 node01 cl_runtime: NOTICE: clcomm: Path node01:net1 -
node02:net2 errors during initiation
May 21 09:42:14 node01 cl_runtime: NOTICE: clcomm: Path node01:net2 -
node02:net3 errors during initiation
```

For security purposes, you must configure Oracle Solaris Cluster to use the Solaris Firewall service.

Workaround: Unblock the Oracle Solaris Cluster interconnect traffic.

Use the appropriate Solaris Firewall service configuration file that is associated with your environment.

- **Oracle Solaris global cluster.** Update the following example Oracle Solaris global cluster `fp.conf` configuration file to replace example interconnect network objects with the objects in your environment.

The example interconnect network objects are `net0`, `net1`, `net2`, `net3`, and `clprivnet0`.

```
# cat /etc/firewall/pf.conf
#!/usr/sbin/pfctl -f

ext_if = "aggr0"
client_out = "{22, 111, 8059, 8060, 8061, 8062, 6499, 11161, 11162, 11163,
11164, 11165}"

# do not filter local interface
set skip on lo0
# do not filter global Cluster interconnects
set skip on net0
set skip on net1
set skip on net2
set skip on net3
set skip on clprivnet0
```

```
# do not filter local ILOM interface
set skip on sp-phys0

block in log quick on egress proto tcp to port { 22 }
block return log all

pass in log proto tcp from any to any port 22 <> 23
pass out log proto tcp from any to any
pass in log proto udp from any to any
pass out inet proto icmp all icmp-type echoreq keep state
pass in log proto icmp from any to any

table <accesslist> persist file "/etc/accesslist"
pass in log on $ext_if proto tcp from <accesslist> to port $client_out flags
S/SA keep state
pass out on $ext_if proto udp all

# Solaris Cluster: allow all connections initiated from this system
pass out

set debug info
```

- **Oracle Solaris exclusive-IP zone cluster.** Update the following example Oracle Solaris exclusive-IP zone cluster `fp.conf` configuration file to replace example interconnect network objects with the objects in your environment.

The example interconnect network objects are `net0_zcdb1_vnic0`, `net1_zcdb1_vnic0`, `net2_zcdb1_vnic0`, `net3_zcdb1_vnic0`, and `clprivnet1`.

```
# cat /etc/firewall/pf.conf
set reassemble yes no-df

ext_if = "zcdb1vnic1"
client_out = "{22, 5201, 111, 8059, 8060, 8061, 8062, 6499, 11161, 11162,
11163, 11164, 11165}"

set skip on lo0

set skip on net0_zcdb1_vnic0
set skip on net1_zcdb1_vnic0
set skip on net2_zcdb1_vnic0
set skip on net3_zcdb1_vnic0
set skip on clprivnet1

block in log quick on egress proto tcp to port { 22 }
block return log all

pass in log proto tcp from any to any port 22 <> 23
pass out log proto tcp from any to any
pass in log proto udp from any to any
pass out inet proto icmp all icmp-type echoreq keep state
pass in log proto icmp from any to any

table <accesslist> persist file "/etc/accesslist"
pass in log on $ext_if proto tcp from <accesslist> to port $client_out flags
S/SA keep state
pass out on $ext_if proto udp all

pass out

set debug error
```

- **Oracle Solaris shared-IP zone cluster.** Unblock the `clprivnet` interconnect network object because a shared-IP zone cluster shares the private interconnects with the global zone.

3

Installation Notes

This chapter contains known issues and bugs that affect the installation of Oracle Solaris Cluster 4.4. Contact Oracle support services to find out whether a code fix has become available.

Installation Issues

`rpc-authdes` Utilities Moved From `core-os` To `legacy/security/rpc-authdes` (30680102)

Problem Summary: Required utilities to set up DES authentication not found on the system by default.

Workaround: To setup DES authentication, install `legacy/security/rpc-authdes` on all nodes of the cluster to obtain required utilities.

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
- nlsps
- 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)

4

Administration Notes

This chapter contains known issues and bugs that affect the administration of Oracle Solaris Cluster 4.4. Contact Oracle support services to find out whether a code fix has become available.

Administration Issues

Depending on `vfstab` Creation of `nfs` rs Fails w INTERNAL ERROR if Failover zpool (28546356)

If `/etc/vfstab` does not have uncommented entries that define a mount point, and the `nfs` resource depends on an `HASStoragePlus` resource for a failover zpool, creation of `nfs` resource fails with "INTERNAL ERROR: Failed to build vfstab entries".

Workaround: Update `/etc/vfstab` so that it has an uncommented entry that defines a mount point.

`cacoadm start` Fails With Warning: found remaining processes in SMF contract (28508074)

Problem Summary: Sometimes booting Oracle Solaris 11.4 and Oracle Solaris Cluster 4.4 after an update results in the `cacoadm start` operation failing and issuing the Warning: found remaining processes in SMF contract warning.

Command output similar to the following indicates that the `cacao` service is not running.

```
# svcs -x svc:/system/cluster/cacao-install:default
State: maintenance since Tue Aug 14 14:41:26 2018
Reason: Start method failed repeatedly, last died on Killed (9).
  See: http://support.oracle.com/msg/SMF-8000-KS
  See: /var/svc/log/system-cluster-cacao-install:default.log
Impact: This service is not running.
```

Workaround: The following commands reset and restart the `cacao` service and verify that the service is running:

```
# svcadm disable cacao-install
# cacoadm prepare-uninstall
# svcadm enable cacao-install
# svcs -a | egrep "cacao|common"
online 08:21:29 svc:/application/management/common-agent-container-1:default
online 08:21:38 svc:/system/cluster/cacao-install:default
online 08:21:38 svc:/system/cluster/cacao-uninstall:default
```

Restrictions for SVMs in Zone Cluster (28384238, 28313846)

At this time, any application that needs to change SVM device permissions in a zone is not supported. Contact an Oracle support representative to learn whether a workaround or fix is available.

Several Cluster Checks Fail With Execution Error (27092918)

Problem Summary: Cluster check S6994589, which checks for Oracle Solaris and Oracle Solaris Cluster release compatibility, and other checks fail with execution error.

Workaround: Contact an Oracle support representative to learn whether a fix is available.

`zoneadm` Commands Hanging at Bootup into 11.4, Cluster 4.4 Following Upgrade (25743399)

Problem Summary: When booting Oracle Solaris 11.4, following an upgrade to Oracle Solaris Cluster 4.4, sometimes `zoneadm` commands hang, immediately after the reboot, when issued against a zone managed by the `ORCL.ha-zone_sczbt` resource type or a zone of a zone cluster.

Workaround: Reboot the node where the commands are hanging.

Workaround: Kill the `zoneadm` command from another terminal window and try again later.

`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 zone 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.

Failed to Remove a File from Global ZFS Filesystem When Quota is Exceeded (20860157)

Problem Summary: You will not be able to remove a file if the quota set on a dataset is exceeded.

Workaround: There are several workarounds for this issue. You can temporarily remove the quota, and then remove the file, or you can truncate one or more file(s) to bring the total filesystem usage under the quota.

One option is to export the zpool and re-import it as a local ZFS file system. Remove the file(s). Re-import the zpool with global mounting.

Another option is to truncate the file. For example, if `file1` is a large file which is causing the filesystem to exceed quota and you want to delete `file1`, you can truncate the file by redirecting a null output into it:

```
schost$ ls -l file1
-rw----- 1 user1      group1      10485760 Oct 21 17:47 file1
schost$ > file1
schost$ ls -l file1
-rw----- 1 user1      group1           0 Oct 21 17:47 file1
schost$ rm file1
```

`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.

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

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

For example:

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

Workaround: Do one of the following:

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

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.

5

Runtime Notes

This chapter contains known issues and bugs that affect the runtime of Oracle Solaris Cluster 4.4. Contact Oracle support services to find out whether a code fix has become available.

Runtime Issues

Investigate Write Performance Improvements for Global ZFS (26201389)

Problem Summary: Based on tested use cases, there might be reasonable performance degradation on PxFS on ZFS compared to PxFS on UFS. Further investigation is needed to address the appropriate usage of Global ZFS.

6

Oracle Solaris Cluster Manager Browser Interface Notes

This chapter contains known issues and bugs that affect the browser interface of Oracle Solaris Cluster 4.4. Contact Oracle support services to find out whether a code fix has become available.

Browser Interface Issues

Selecting or Unselecting Accessibility Item Results in Blank Screen (28003165)

Problem Summary: After logging in to the GUI, then clicking on an accessibility menu item the screen goes blank.

Workaround: To resolve the page view issue and use the accessibility features, follow these steps:

1. Refresh the page and logout of the GUI.
2. Make the desired selection in the accessibility menu.
3. Then log in to the cluster.

On Newly Booted Nodes Sometimes Browser Becomes Unresponsive (26675695)

Problem Summary: After the fresh boot of a system, the first time you try to access the browser interface, it is possible that the browser interface can become unresponsive.

Workaround: Close and reopen the browser.

Tasks - Highly Available Storage Can Not Proceed in ES (23135263)

Problem Summary: When attempting to create high availability storage using the wizard in the Spanish language browser interface, the wizard freezes and the operation fails.

Workaround: Refresh or re-log into the browser interface.

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

Workaround: 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.

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.

RAC Wizard Fails w `java.lang.NullPointerException` When Creating Clusterware rs (27553809)

See Oracle BugDB for workaround information.

Oracle DB GUI Wizard Does Not Show Any ASM Projects Leading to Failure (27548953)

See Oracle BugDB for workaround information.

ADF Error Listing RG's For Add to PG Table (25065856)

See Oracle BugDB for workaround information.

7

Data Services Notes

This chapter contains known issues and bugs that affect data services of Oracle Solaris Cluster 4.4. Contact Oracle support services to find out whether a code fix has become available.

Data Services Issues

Resource Creation Using `SUNW.scalable_rac_server_proxy` fails - Specified `DB_NAME db_name` Not Registered with CRS (29054450)

Problem Summary: There can be a problem using an Oracle E-Business Suite delivered Oracle Home with the `SUNW.scalable_rac_server_proxy` RT to create a resource. That resource creation fails with "Validation failed. Specified `DB_NAME db_name` not registered with CRS" Refer to the bug for more information.

Workaround: Copy the `osdbagrp` binary from other nodes of the cluster if it is working fine on the other nodes.

Migration of Resources Registered as `ORCL.ha-zone_sczbt`, `ORCL.ha-zone_sczsh` Or `ORCL.ha-zone_sczsmf` to the Latest Resource Type Version (21926061)

Problem Summary: For resources registered as `ORCL.ha-zone_sczbt`, `ORCL.ha-zone_sczsh` or `ORCL.ha-zone_sczsmf`, before upgrading to Oracle Solaris Cluster 4.4, you must upgrade to Oracle Solaris Cluster 4.3 SRU 4 or later and specifically implement note 14 from the OSC 4.3 SRU4 README (Doc ID 2170110.1).

Note 14: Migration of resources registered as `ORCL.ha-zone_sczbt`, `ORCL.ha-zone_sczsh` or `ORCL.ha-zone_sczsmf` to the latest resource type version (Bug 21926061)

You must upgrade an existing resource of resource type `ORCL.ha-zone_component` to the latest resource type version by using the `/opt/SUNWsczone/component/util/rt_upgrade` script.

Warning: If performing a rolling upgrade, you must wait until all cluster nodes have been upgraded before running the `/opt/SUNWscxvm/util/rt_upgrade` script.

Example 7-1 Upgrading a single resource of resource type `ORCL.ha-zone_component`

```
# /opt/SUNWsczone/component/util/rt_upgrade resource
```

Example 7-2 Upgrading multiple resources of resource type `ORCL.ha-zone_component`

```
# /opt/SUNWsczone/util/component/rt_upgrade resource1 resource2
```

Example 7-3 Upgrading all existing resources of resource type ORCL.ha-zone_component

```
# /opt/SUNWscxvm/util/component/rt_upgrade +
```

where *component* is *sczbt*, *sczsh*, or *sczsmf* for ORCL.ha-zone_sczbt, ORCL.ha-zone_sczsh, and ORCL.ha-zone_sczsmf resource type respectively.

Once upgraded to Oracle Solaris Cluster 4.4, resources registered as ORCL.ha-zone_sczbt, ORCL.ha-zone_sczsh or ORCL.ha-zone_sczsmf need to be upgraded again to a newer version, e.g.

```
/usr/cluster/bin/clrt register ORCL.ha-zone_scbt
/usr/cluster/bin/clrs unmonitor <resource>
/usr/cluster/bin/clrs set -p type_version=5 <resource>
/usr/cluster/bin/clrt unregister ORCL.ha-zone_scbt:4
```

 **Note:**

If required, repeat the steps above to upgrade to newer versions of ORCL.ha-zone_sczsh or ORCL.ha-zone_sczsmf.

For more information, refer to [Upgrading a Resource Type in Planning and Administering Data Services for Oracle Solaris Cluster 4.4](#).

HA-LDOMs Resources Must Update to 4.3 before Updating to 4.4 (23308155)

For resources registered as SUNW.ldom, before upgrading to Oracle Solaris Cluster 4.4, you must upgrade to Oracle Solaris Cluster 4.3 SRU 4 or later and specifically implement note 17 from the Oracle Solaris Cluster 4.3 SRU4 README ([Doc ID 2170110.1](#)).

Note 17: Migration of resources registered as SUNW.ldom with RT version up to 8 to the latest resource type version (Bug 23308155).

You must upgrade an existing resource to the latest resource type version by using the `/opt/SUNWscxvm/util/rt_upgrade` script.

Warning: If performing a rolling upgrade, you must wait until all cluster nodes have been upgraded before running the `/opt/SUNWscxvm/util/rt_upgrade` script.

Example 7-4 Upgrading a single SUNW.ldom resource

```
# opt/SUNWscxvm/util/rt_upgrade ldom-rs1 ldom-rs2
```

Example 7-5 Upgrading multiple SUNW.ldom resources

```
# /opt/SUNWscxvm/util/rt_upgrade ldom-rs
```

Example 7-6 Upgrading all existing SUNW.ldom resources

```
# /opt/SUNWscxvm/util/rt_upgrade +
```

Warning: If performing a rolling upgrade with Ldoms using live migration, you must unmonitor the Ldom resources:

```
# clrs unmonitor <ldom-rs>
```

Once upgraded to Oracle Solaris Cluster 4.4, resources registered as `SUNW.ldom` need to be upgraded again to a newer version, for example:

```
/usr/cluster/bin/clrt register SUNW.ldom  
/usr/cluster/bin/clrs unmonitor <resource>  
/usr/cluster/bin/clrs set -p type_version=11 <resource>  
/usr/cluster/bin/clrt unregister SUNW.ldom:10
```

For more information, refer to [Upgrading a Resource Type in Planning and Administering Data Services for Oracle Solaris Cluster 4.4](#).

run_setuid_prog Failing With ORA-01017: Invalid username/password; Logon Denied (21509291)

Problem Summary: Under certain conditions a resource of type `SUNW.oracle_server` fails to start the Oracle database. By default, `SUNW.oracle_server` uses an internal program `run_setuid_prog` that fails if the file group owner is different between `$ORACLE_HOME/bin/oracle` and `$ORACLE_HOME/bin/sqlplus`.

Workaround: On each cluster node where a resource of type `SUNW.oracle_server` will execute, issue the following command:

```
# touch /opt/SUNWscor/.use_su
```

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).

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 filesystem 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.

8

Disaster Recovery Framework Notes

This chapter contains known issues and bugs affect the Oracle Solaris Cluster 4.4 software. Contact Oracle support services to find out whether a code fix has become available.

Disaster Recovery Framework Issues

Update How to Perform a Failback-Switchover on a System That Uses Oracle ZFS Storage Appliance Replication Procedure (36175231)

The following updated procedure replaces [How to Perform a Failback-Switchover on a System That Uses Oracle ZFS Storage Appliance Replication in Oracle Solaris Cluster Remote Replication Guide for Oracle ZFS Storage Appliance](#).

This procedure refers to the following cluster names and ZFS storage appliance names:

- **Original Primary Cluster:** `cluster-paris` and the associated appliance name is `paris`.
- **Current Primary Cluster:** `cluster-newyork` and the associated appliance name is `newyork`.

Use this procedure to restart an application on the original primary cluster, `cluster-paris`, after the data on this cluster has been resynchronized with the data on the current primary cluster, `cluster-newyork`.



Note:

The failback procedures apply only to clusters in a partnership. You need to perform the following procedure only once per partnership.

Before You Begin

Before you perform a failback-switchover, a takeover has occurred on `cluster-newyork`. Ensure that the clusters have the following roles:

- If the original primary cluster had been down, the cluster has been booted and that the disaster recovery framework is enabled on the cluster. For more information about booting a cluster, see [Booting a Cluster in Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4](#).
 - The protection group on the current primary cluster has the primary role.
 - The protection group on the original primary cluster has either the primary role or secondary role, depending on whether the original primary cluster can be reached during the takeover from the current primary cluster.
1. Synchronize the replication from the `newyork` appliance to the `paris` appliance.

You must perform this task to complete a recovery when the cluster experiences a complete site failure or a takeover. Note that the data stores at `cluster-newyork` will have changed and must be replicated on `cluster-paris` before it can return to service.

Perform these steps for each replicated project in the protection group.

- a. Access the Oracle ZFS Storage Appliance browser user interface (BUI) on the `cluster-newyork` site.
- b. Select the project that you want to replicate from the Shares > Projects menu.
- c. Select Replication for the project and the click Update or Sync Now.

If Update Now or Sync Now do not appear as options, click Enable Action to start the replication. Ensure that the replication mode is set to continuous. If neither syncing a replication nor enabling a replication succeed, perform the following steps to execute a manual replication to synchronize the two sites:

- Remove the replication action
- Re-create the replication action
- Enable the replication action in continuous mode

2. Ensure that the protection group is stopped at the `cluster-paris` site.

- a. Determine whether the protection group on the original primary cluster, `cluster-paris`, is active.

```
phys-paris-1# geoadm status
```

- b. Stop the protection group on the original primary cluster if active.

```
phys-paris-1# geopg stop -e local protection-group
```

protection-group specifies the name of the protection group.

- c. Verify that the protection group is stopped.

```
phys-paris-1# geoadm status
```

3. Remove any obsolete projects from the appliance at the `cluster-paris` site.

First, access the BUI on the `cluster-paris` site. Then, from the Shares > Projects menu, manually delete any projects that are listed in the protection group.

4. Resynchronize the original primary cluster, `cluster-paris`, with the current primary cluster, `cluster-newyork`.

The `cluster-paris` cluster forfeits its own configuration and replicates the `cluster-newyork` configuration locally. Then, resynchronize both the partnership configuration and the protection group configuration.

- a. On `cluster-paris`, resynchronize the partnership.

```
phys-paris-1# geops update partnership
```

partnership specifies the name of the partnership.

 **Note:**

Perform this step only one time per partnership, even if you are performing a failback-switchover for multiple protection groups in the partnership.

For more information about synchronizing partnerships, see [Resynchronizing a Partnership in Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4](#).

- b. On `cluster-paris`, resynchronize each protection group.

Because the local role of the protection group on `cluster-newyork` is now primary, this step ensures that the role of the protection group on `cluster-paris` becomes secondary.

```
phys-paris-1# geopg update protection-group
```

For more information about synchronizing protection groups, see [Resynchronizing a Protection Group in Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4](#).

5. On `cluster-paris`, validate the cluster configuration for each protection group.

Ensure that the protection group is not in an error state or the protection group cannot be started.

```
phys-paris-1# geopg validate protection-group
```

protection-group specifies a unique name that identifies a single protection group.

For more information, see [Validating a Protection Group in Installing and Configuring the Disaster Recovery Framework for Oracle Solaris Cluster 4.4](#).

6. On `cluster-paris`, activate each protection group.

Because the protection group on `cluster-paris` has a role of secondary, the `geopg start` command does not restart the application on `cluster-paris`.

```
phys-paris-1# geopg start -e global protection-group
```

The `-e` option specifies the scope of the command, where `local` relates to the local cluster and where `global` relates to both clusters. *protection-group* specifies the name of the protection group.

 **Note:**

Do not use the `-n` option when performing a failback-switchover. You must synchronize from the current primary, `cluster-newyork`, to the current secondary, `cluster-paris`.

Because the protection group fills a secondary role, the data is synchronized from the current primary, `cluster-newyork`, to the current secondary, `cluster-paris`.

For more information about the `geopg start` command, see [How to Activate a Protection Group in Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4](#).

7. Confirm that the data is completely synchronized.

The data is completely synchronized when the state of the protection group on `cluster-newyork` is OK. The protection group has a local state of OK when the appliance data store on `cluster-newyork` is being updated to the `cluster-paris` cluster.

To confirm that the state of the protection group on `cluster-newyork` is OK, view the Protection Group section in the `geoadm status` command output.

8. On both partner clusters, ensure that the protection group is activated.

```
# geoadm status
```

9. On either cluster, perform a switchover from `cluster-newyork` to `cluster-paris` for each protection group.

```
# geopg switchover [-f] -m cluster-paris protection-group
```

For more information, see [How to Switch Over Replication From the Primary Cluster to the Secondary Cluster in Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4](#).

`cluster-paris` resumes its original role as primary cluster for the protection group.

10. Ensure that the switchover is successful.

Verify that the protection group is now primary on `cluster-paris` and is secondary on `cluster-newyork`. Also verify that the state of Data Replication and Resource Groups is OK on both clusters.

```
# geoadm status
```

Check the runtime status of the application resource group and replication for each protection group in the Status field and Status Message field for the remote replication component.

```
# clresourcegroup status -v protection-group
```

ZFS SA Replication Module Does Not Support Cluster File Systems With ZFS (29223365)

Problem Summary: Currently cluster file systems with ZFS where zpool LUNs are being replicated using ZFS SA replication is not supported by the disaster recovery framework. This is due to the HASP resource leaving the zpool imported at the former primary cluster after the protection group is switched to a new primary cluster.

Workaround: Use ZFS snapshot replication with cluster filesystems with ZFS.

Case Where Node Stays Fenced Off & Later `ScalMountPoint rs` Fails to Start There (28362456)

Problem Summary: When a cluster node goes down, it stays fenced off from the NFS share, and later `ScalMountPoint rs` fails to start there when the protection group is switched back to this node.

Workaround: Set the share mode to `RW` for this node, then restart the `ScalMountPoint` `rg` on the node where it failed to start.

Flag Cases & Disallow Switchovers Accordingly Where Dependent `rgs` are Not in `pg` (28004581)

Problem Summary: If there are any `RG_affinities` like `++/+++` or `RG_dependencies` from Resource group(s) which are outside Protection group on to the Resource group(s) which is in the Protection group, and if those external Resource group(s) are not `unmanaged`, then switchover of Protection group fails.

Workaround: If you must keep dependent resource groups outside of the protection group managing the dependee resource group(s), consider the following options:

- Put those resource groups in other protection group(s), then put these protection groups in the same multigroup as the dependee protection group with dependency from the protection group that contains the dependent resource group(s) on this dependee protection group. To initiate a switchover, use the `geomg` command with the multigroup and not `geopg` with the individual protection groups.
- Still keep the dependent resource groups outside of protection groups, but you must bring them to the `unmanaged` state before initiating the switchover of the protection groups containing the dependee resource groups.

`geopg` Returns Java Exception on 4.4 ZC Node While Other Node is Running 4.3.9 (27148373)

Problem Summary: After a rolling update to Oracle Solaris Cluster 4.4 from Oracle Solaris Cluster 4.3, `geopg list` and `geopg show` fails with a Java exception.

Workaround: Switch the resource `geo-failovercontrol` from the cluster node running Oracle Solaris Cluster 4.3 on to the node which runs Oracle Solaris Cluster 4.4.

Switchover Fails if Replicated FS Exported & Mountpoint Can be Overwritten (25970289)

Problem Summary: A protection group switchover fails if replicated filesystems from a project in the protection group are exported. Also, if such a filesystem is exported and its mountpoint changed, this results in mountpoint getting overwritten.

Workaround: If you must export replicated filesystems, you must stop exporting them before a switchover. In addition, if the mountpoint of the exported replicated filesystems is changed, you must set it to the same mountpoint of the source before the switchover.

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`.

Bug 16964103 About Underlying `did` Not Available Still Exists if `pg` in ZC (24582738)

Problem Summary: In some cases, the ASM device group is unable to come online at switchover because the underlying device is not accepting I/O.

Workaround: To restart the application resource group:

1. Issue `cldevice populate` from one node of global cluster.
2. From one node of the zone cluster, restart the application resource group that failed due to no access to storage (for example, `clrg restart <rg>`).

Dataguard Manager User for Container `db (c##dgmgr)` Fails (24354703)

Problem Summary: When creating a data guard protection group with the dataguard manager user set to `sysdba_username=c##dgmgr`, the operation fails with the following message:

```
Unable to connect using connect string "c##dgmgr/<password>@<instance>".  
Check that the password or connect string is correct and that the database instance is  
running
```

Workaround: To solve this issue, set `sysdba_username=sys`.

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 (22529981)

Problem Summary: Projects or mount points configured with the same name on the target appliance as the one on source appliance managed by Disaster Recovery Framework 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.

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.

Disaster Recovery Framework Does Not Support ZFSSA Offline Replication Feature (21934145)

Problem Summary: Disaster Recovery Framework 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 Disaster Recovery Framework.

9

Documentation Notes

This chapter discusses errors or omissions for documentation in the Oracle Solaris Cluster 4.4 release.

`clresource create -d` [Commands Incorrectly Specify Resource Names in Oracle Solaris Cluster Data Service for Oracle E-Business Suite as of Release 12.2 Guide \(26401103\)](#)

The incorrect resource names are missing a hyphen between `-resource` and `name`. The correct form of the resource name is `*-resource-name`, not `*-resource name`.

These incorrect resource names appear in the `clresource create` command synopses and in the resource definitions of the following tasks:

- [How to Create the Root Service Group Resource in Oracle Solaris Cluster Data Service for Oracle E-Business Suite as of Release 12.2 Guide](#)

Step 2: Replace the `clresource create` command synopsis with the following:

```
# clresource create -g resource group -t ORCL.ebs \  
-p base_dir=apps base dir \  
-p service_group=root -p service=node_manager \  
-p resource_dependencies_offline_restart=storage-resource \  
[ -p interpose_logical_hostname=logical host ] \  
[ -p resource_dependencies=logical-host-resource ] \  
-d node-manager-resource-name
```

Replace the resource definition with the following:

`node-manager-resource-name`

Specifies an Oracle Solaris Cluster resource name for the Node Manager. This name must be unique with Oracle Solaris Cluster.

- [How to Create the Web Administration Service Group Resources in Oracle Solaris Cluster Data Service for Oracle E-Business Suite as of Release 12.2 Guide](#)

– **Step 2:** Replace the `clresource create` command synopsis with the following:

```
# clresource create -g resource group -t ORCL.ebs \  
-p base_dir=apps base dir \  
-p service_group=web_admin -p service=tns_apps \  
-p resource_dependencies_offline_restart=storage-resource \  
-p pmf_managed=true \  
-p interpose_logical_hostname=logical host \  
-p resource_dependencies=logical-host-resource \  
-d tns-listener-resource-name
```

Replace the resource definition with the following:

tns-listener-resource-name

Specifies an Oracle Solaris Cluster resource name for the TNS Listener. This name must be unique with Oracle Solaris Cluster.

- **Step 3:** Replace the `clresource create` command synopsis with the following:

```
# clresource create -g resource group -t ORCL.ebs \
-p base_dir=apps base dir \
-p service_group=web_admin -p service=admin_server \
-p resource_dependencies_offline_restart=storage-resource,tns-listener-
resource,node-manager-resource,[scope:]db-resource \
-p interpose_logical_hostname=logical host \
-p resource_dependencies=logical-host-resource \
-d admin-server-resource-name
```

Replace the resource definition with the following:

admin-server-resource-name

Specifies an Oracle Solaris Cluster resource name for the WebLogic Admin Server. This name must be unique with Oracle Solaris Cluster.

- [How to Create the Web Entry Point Service Group Resources in Oracle Solaris Cluster Data Service for Oracle E-Business Suite as of Release 12.2 Guide](#)

- **Step 2:** Replace the `clresource create` command synopsis with the following:

```
# clresource create -g resource group -t ORCL.ebs \
-p base_dir=apps base dir \
-p service_group=web_entry-p service=opmn \
-p resource_dependencies_offline_restart=storage-resource \
-p pmf_managed=true \
[ -p interpose_logical_hostname=logical host ] \
[ -p resource_dependencies=logical-host-resource ] \
-d opmn-resource-name
```

Replace the resource definition with the following:

opmn-resource-name

Specifies an Oracle Solaris Cluster resource name for the OPMN Server. This name must be unique with Oracle Solaris Cluster.

- **Step 3:** Replace the `clresource create` command synopsis with the following:

```
# clresource create -g resource group -t ORCL.ebs \
-p base_dir=apps base dir \
-p service_group=web_entry-p service=ohs \
-p resource_dependencies_offline_restart=storage-resource \
-p resource_dependencies=opmn-resource,logical-host-resource \
[ -p interpose_logical_hostname=logical host ] \
-d ohs-resource-name
```

Replace the resource definition with the following:

ohs-resource-name

Specifies an Oracle Solaris Cluster resource name for OHS. This name must be unique with Oracle Solaris Cluster.

- [How to Create the Web Application Service Group Resources in Oracle Solaris Cluster Data Service for Oracle E-Business Suite as of Release 12.2 Guide](#)

- **Step 3:** Replace the `clresource create` command synopsis with the following:

```
# clresource create -g resource group -t ORCL.ebs \
-p base_dir=apps base dir \
-p service_group=web_applications -p service=tns_apps \
-p resource_dependencies_offline_restart=storage-resource \
-p pmf_managed=true \
[ -p interpose_logical_hostname=logical host ] \
[ -p resource_dependencies=logical-host-resource ] \
-d tns-listener-resource-name
```

Replace the resource definition with the following:

tns-listener-resource-name

Specifies an Oracle Solaris Cluster resource name for the TNS Listener. This name must be unique with Oracle Solaris Cluster.

- **Step 4:** Replace the `clresource create` command synopsis with the following:

```
# clresource create -g resource group -t ORCL.ebs \
-p base_dir=apps base dir \
-p service_group=web_applications -p service=oaocore \
-p resource_dependencies=tns-listener-resource,node-manager-resource,
[scope:]db-resource \
-p resource_dependencies_offline_restart=storage-resource \
[ -p interpose_logical_hostname=logical host ] \
[ -p resource_dependencies=logical-host-resource ] \
-d oaocore-resource-name
```

Replace the resource definition with the following:

oaocore-resource-name

Specify an Oracle Solaris Cluster resource name for the `oaocore` Web Application. This name must be unique with Oracle Solaris Cluster.

- **Step 5:** Replace the `clresource create` command synopsis with the following:

```
# clresource create -g resource group -t ORCL.ebs \
-p base_dir=apps base dir \
-p service_group=web_applications -p service=oafm \
-p resource_dependencies=tns-listener-resource,node-manager-resource,
[scope:]db-resource \
-p resource_dependencies_offline_restart=storage-resource \
[ -p interpose_logical_hostname=logical host ] \
[ -p resource_dependencies=logical-host-resource ] \
-d oafm-resource-name
```

Replace the resource definition with the following:

oafm-resource-name

Specifies an Oracle Solaris Cluster resource name for the Web Application `oafm`. This name must be unique with Oracle Solaris Cluster.

- **Step 6:** Replace the `clresource create` command synopsis with the following:

```
# clresource create -g resource group -t ORCL.ebs \
-p base_dir=apps base dir \
-p service_group=web_applications -p service=forms \
-p resource_dependencies=tns-listener-resource,node-manager-resource,
```

```
[scope:]db-resource \  
-p resource_dependencies_offline_restart=storage-resource \  
[ -p interpose_logical_hostname=logical host ] \  
[ -p resource_dependencies=logical-host-resource ] \  
-d forms-resource-name
```

Replace the resource definition with the following:

forms-resource-name

Specify an Oracle Solaris Cluster resource name for the `forms` web application. This name must be unique with Oracle Solaris Cluster.

- **Step 7:** Replace the `clresource create` command synopsis with the following:

```
# clresource create -g resource group -t ORCL.ebs \  
-p base_dir=apps base dir \  
-p service_group=web_applications -p service=forms-c4ws \  
-p resource_dependencies=tns-listener-resource,node-manager-resource,  
[scope:]db-resource \  
-p resource_dependencies_offline_restart=storage-resource \  
[ -p interpose_logical_hostname=logical host ] \  
[ -p resource_dependencies=logical-host-resource ] \  
-d forms-c4ws-resource-name
```

Replace the resource definition with the following:

forms-c4ws-resource-name

Specify an Oracle Solaris Cluster resource name for the `forms-c4ws` web application. This name must be unique with Oracle Solaris Cluster.

- [How to Create the Batch Service Group Resources in Oracle Solaris Cluster Data Service for Oracle E-Business Suite as of Release 12.2 Guide](#)

- **Step 2:** Replace the `clresource create` command synopsis with the following:

```
# clresource create -g resource group -t ORCL.ebs \  
-p base_dir=apps base dir \  
-p service_group=batch -p service=tns_apps \  
-p resource_dependencies_offline_restart=storage-resource \  
-p pmf_managed=true \  
[ -p interpose_logical_hostname=logical host ] \  
[ -p resource_dependencies=logical-host-resource ] \  
-d tns-listener-resource-name
```

Replace the resource definition with the following:

tns-listener-resource-name

Specifies an Oracle Solaris Cluster resource name for the TNS Listener. This name must be unique with Oracle Solaris Cluster.

- **Step 3:** Replace the `clresource create` command synopsis with the following:

```
# clresource create -g resource group -t ORCL.ebs \  
-p base_dir=apps base dir \  
-p service_group=batch -p service=concmgr \  
-p resource_dependencies=tns-listener-resource,[scope:]db-resource \  
-p resource_dependencies_offline_restart=storage-resource \  
[ -p interpose_logical_hostname=logical host ] \  
\
```



```
[ -p resource_dependencies=logical-host-resource ] \  
-d concmgr-resource-name
```

Replace the resource definition with the following:

concmgr-resource-name

Specifies an Oracle Solaris Cluster resource name for the Concurrent manager. This name must be unique with Oracle Solaris Cluster.

- [How to Create the Other Service Group Resource in Oracle Solaris Cluster Data Service for Oracle E-Business Suite as of Release 12.2 Guide](#)

Step 2: Replace the `clresource create` command synopsis with the following:

```
# clresource create -g resource group -t ORCL.ebs \  
-p base_dir=apps base dir \  
-p service_group=other -p service=forms_server \  
-p resource_dependencies_offline_restart=storage-resource \  
-p pmf_managed=true \  
[ -p interpose_logical_hostname=logical host ] \  
[ -p resource_dependencies=logical-host-resource ] \  
-d forms-srv-resource-name
```

Replace the resource definition with the following:

forms-srv-resource-name

Specifies an Oracle Solaris Cluster resource name for the Forms Server (Socket mode). This name must be unique with Oracle Solaris Cluster.

Refer to Oracle VM Server for SPARC 3.6 Documentation in Oracle Solaris Cluster Data Service for Oracle VM Server for SPARC Guide (33890482)

Step 3 of SPARC: How to Configure HA for Oracle VM Server in *Oracle Solaris Cluster Data Service for Oracle VM Server for SPARC Guide* incorrectly refers to the Oracle VM Server for SPARC 3.2 documentation.

Replace this reference with [Configuring SSL Certificates for Migration in Oracle VM Server for SPARC 3.6 Administration Guide](#).

`clpstring` Command in Oracle Solaris Cluster Data Service for Oracle VM Server for SPARC Guide Has Incorrect Parameters (33837717)

Step 3 of [SPARC: How to Configure HA for Oracle VM Server for SPARC in Oracle Solaris Cluster Data Service for Oracle VM Server for SPARC Guide](#) incorrectly specifies the `-t` option in the `clstring` command.

Replace the following command:

```
# /usr/cluster/bin/clpstring create -b rname \  
-t rname_rname_ldompasswd
```

Enter string value:
Enter string value again:

With the following command:

```
# /usr/cluster/bin/clpstring create -b rname rgname_rname_ldompasswd
Enter string value:
Enter string value again:
```

Note that you must specify the root password.

Add `parameter_file` Property Section to Setting HA for Oracle Database Extension Properties in *Oracle Solaris Cluster Data Service for Oracle Database Guide (33131714)*

The `parameter_file` property specifies the location of an ASCII parameter file (pfile). The file you specify can also include the location of a binary version of the pfile called an spfile.

When you specify a `parameter_file` property value, the startup command run by the `SUNW.oracle_server` resource includes `pfile=location-specified-by-$PARAMETER_FILE`.

If you want to use an spfile instead of a pfile, use one of the following options:

- Specify an empty value for the `parameter_file` property so that the database uses its default search path to find the startup command.
- Create a pfile that contains the full path to a binary spfile. Then, specify the full path to the pfile as the value of the `parameter_file` property. The pfile references the full path to the spfile as a value of the `SPFILE` variable. For example, the `$ORACLE_HOME/dbs/init$ORACLE_SID.ora` pfile includes a line similar to `SPFILE=$ORACLE_HOME/dbs/spfile$ORACLE_SID.ora`. Then, update the Oracle Solaris Cluster HA Oracle resource's `parameter_file` property to specify this pfile.
- Use the spfile to create a pfile and then reference the new pfile as the `parameter_file` property value. At the `SQL>` prompt, specify the `CREATE PFILE = 'init$ORACLE_SID.ora' FROM SPFILE = 'spfile$ORACLE_SID.ora';` command. Then, update the Oracle Solaris Cluster HA Oracle resource's `parameter_file` property to specify this pfile.

Update Instructions About Booting All Nodes Into Non-Cluster Mode and Rebooting Nodes in *Administering an Oracle Solaris Cluster 4.4 Configuration (31092326)*

Replace Step 1 and Step 9 of [How to Change the Private Network Address or Address Range of an Existing Cluster in *Administering an Oracle Solaris Cluster 4.4 Configuration*](#) with the following content:

Step 1:

1. Reboot all cluster nodes into non-cluster mode by performing the following substeps:

a. On one node of the cluster, assume the root role or a role that provides `solaris.cluster.admin` authorization.

b. Shut down the cluster.

```
# cluster shutdown -g0 -y
```

c. On each node of the cluster, boot the node into non-cluster mode.

Run one of the following commands:

- On SPARC, run `boot -x` at OK prompt.

- On x86, specify the `-x` option in GRUB.

Step 9:

9. Reboot each cluster node back into cluster mode.

```
# reboot
```

Update Database Preparation in *Oracle Solaris Cluster Data Service for Oracle GoldenGate Guide* (20170637)

Replace the final `mysql` command of [Database Preparation in Oracle Solaris Cluster Data Service for Oracle GoldenGate Guide](#) with the following command:

```
root@vzsomala:~# /usr/local/mysql/bin/mysql -h soma-1 -ugg -pgg
mysql> use golden
mysql> create table tab1 (row1 int not null, row2 char(64));
mysql> create unique index ind1 on tab1(row1);
mysql> exit
```

Add How to Add a ZFS Cluster File System to a Zone Cluster (CLI) Task to Adding File Systems to a Zone Cluster in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4* (30930843)

Include the following reference to the Adding File Systems to a Zone Cluster task list: [How to Add a ZFS Cluster File System to a Zone Cluster \(CLI\) in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*](#).

Change Instances of StorageTek to HSM in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4* (30930843)

Throughout this book, replace references to StorageTek QFS with HSM.

Add Clarification About a SAP xserver to How to Register and Configure an SAP xserver Resource in *Oracle Solaris Cluster Data Service for SAP MaxDB Guide* (21803412)

HA for SAP MaxDB requires that a SAP xserver starts on any node to which the SAP MaxDB resource can fail over. For this reason, a SAP xserver must be configured on a global file system.

Oracle Database on Oracle Solaris Cluster Does Not Require oracle:dba User and Group Settings (32988568)

Oracle Solaris Cluster Data Service for Oracle Database Guide shows that the Oracle Database uses the `oracle:dba` user and group settings. However, the agent does not require these values. For information about the user and group settings that you can use, see your database documentation.

`SUNW.oracle_server` (7) Man Page Missing From *Oracle Solaris Cluster 4.4 Reference Manual* (31191927)

While the `SUNW.oracle_server(7)` man page is not included in *Oracle Solaris Cluster 4.4 Reference Manual*, you can access the man page by running the following `man` command:

```
$ man SUNW.oracle_server
```

Clarify That How to Remove a Node From a Zone Cluster in *Administering an Oracle Solaris Cluster 4.4 Configuration* Removes the Zone From the Cluster Node (31145133)

The [How to Remove a Node From a Zone Cluster](#) task not only removes the node from the zone cluster, it also removes the zone from the physical node.

- Replace the first sentence of the first paragraph with the following paragraph:
You can remove a node from a zone cluster by halting the node, uninstalling it, and removing the node from the configuration. This procedure removes the zone from the cluster node.
- In Step 5 of [How to Remove a Node From a Zone Cluster in Administering an Oracle Solaris Cluster 4.4 Configuration](#), add the following note before showing the `clzonecluster` command:

 **Note:**

This step removes the *zone-cluster-name* zone from the *node* physical-node property value specified in the remove command.

Incorrect Configuration File Name in Oracle Solaris Cluster Data Service for Apache Tomcat Guide (32973775)

Step 5 of [How to Activate Debugging for HA for Apache Tomcat in Oracle Solaris Cluster Data Service for Apache Tomcat Guide](#) mistakenly identifies the configuration file as `/opt/SUNWsczone/sczbt/etc/config`. The correct configuration file name is `/opt/SUNWsctomcat/etc/config`.

Update Erroneous `clpstring` Command in Oracle Solaris Cluster Data Service for Oracle VM Server for SPARC Guide (32988568)

Replace the third list item in Step 3 of [SPARC: How to Configure HA for Oracle VM Server in Oracle Solaris Cluster Data Service for Oracle VM Server for SPARC Guide](#) with the following:

- **For the resource type version 6 or later, create a private string named `rgname_rname_ldompasswd` to store the target host password that is required for guest domain migration.**

Type the following command on any one node in the cluster:

```
# /usr/cluster/bin/clpstring create -b rname -t resource rgname_rname_ldompasswd
Enter string value:
Enter string value again:
```

The `cplstring create` command creates the `SUNW.ldom` resource that has `rname` as its resource name and `rgname` as its resource group.

Update Misspelled Property Name in Oracle Solaris Cluster Data Service for Oracle VM Server for SPARC Guide (32988568)

Step 4 of [SPARC: How to Configure HA for Oracle VM Server in Oracle Solaris Cluster Data Service for Oracle VM Server for SPARC Guide](#) includes three misspelled instances of the `Migration_type` property name. Replace all instances of `Miration_type` with `Migration_type`.

List of Zone Cluster Tasks (31080689)

This section includes a complete list of tasks that are related to managing zone clusters. Note that some tasks appear in more than one Oracle Solaris Cluster book.

- [How to Create a Zone Cluster \(clsetup\) in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)
- [How to Create a solaris10 Brand Zone Cluster \(CLI\) in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*](#)
- [Creating the Image for Directly Migrating Oracle Solaris 10 Systems Into Zones in *Creating and Using Oracle Solaris 10 Zones*](#)
- [How to Install a Zone Cluster from the Unified Archive in *Administering an Oracle Solaris Cluster 4.4 Configuration*](#)
- [How to Install a Zone Cluster from the Unified Archive in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)
- [How to Configure a Zone Cluster from the Unified Archive in *Administering an Oracle Solaris Cluster 4.4 Configuration*](#)
- [How to Configure a Zone Cluster from the Unified Archive in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)
- [How to Configure a Zone Cluster to Use Trusted Extensions in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)
- [How to Add a Node to an Existing Cluster or Zone Cluster in *Administering an Oracle Solaris Cluster 4.4 Configuration*](#)
- [How to Add a Network Address to a Zone Cluster in *Administering an Oracle Solaris Cluster 4.4 Configuration*](#)
- [How to Add a Network Address to a Zone Cluster in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)
- [How to Add an HSM \(StorageTek QFS\) Shared File System to a Zone Cluster \(CLI\) in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)
- [How to Add a Cluster File System to a Zone Cluster \(clsetup\) in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)
- [How to Add a DID Device to a Zone Cluster \(CLI\) in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)
- [How to Add a Disk Set to a Zone Cluster \(CLI\) in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)
- [How to Add a Global Storage Device to a Zone Cluster \(clsetup\) in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)
- [How to Add a Highly Available Local File System to a Zone Cluster \(clsetup\) in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)
- [How to Add a Highly Available Local File System to a Zone Cluster \(CLI\) in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)
- [How to Add a Local ZFS Storage Pool to a Specific Zone-Cluster Node \(CLI\) in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)

- [How to Add a Raw-Disk Device to a Specific Zone-Cluster Node \(CLI\) in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)
- [How to Add a UFS Cluster File System to a Zone Cluster \(CLI\) in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)
- [How to Add a Local File System to a Specific Zone-Cluster Node \(CLI\) in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)
- [How to Add a ZFS Storage Pool to a Zone Cluster \(CLI\) in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)
- [How to Add a ZFS Storage Pool to a Zone Cluster \(clsetup\) in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)
- [How to Add an Oracle ACFS File System to a Zone Cluster \(CLI\) in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)
- [How to Add a Node to an Existing Cluster or Zone Cluster in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)
- [How to Import an Installed Zone into a New Zone Cluster in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)
- [How to Import an Installed Zone into an existing Zone Cluster in *Creating and Administering Zone Clusters With Oracle Solaris Cluster 4.4*](#)
- [How to Remove a Node From a Zone Cluster in *Administering an Oracle Solaris Cluster 4.4 Configuration*](#)
- [How to Remove a File System From a Zone Cluster in *Administering an Oracle Solaris Cluster 4.4 Configuration*](#)
- [How to Remove a Storage Device From a Zone Cluster in *Administering an Oracle Solaris Cluster 4.4 Configuration*](#)
- [How to Unconfigure a Zone Cluster in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*](#)

Add Requirements to How to Use Link Aggregation Control Protocol (LACP) with Cluster Interconnects in *Administering an Oracle Solaris Cluster 4.4 Configuration (33099699)*

[How to Use Link Aggregation Control Protocol \(LACP\) with Cluster Interconnects in *Administering an Oracle Solaris Cluster 4.4 Configuration*](#) is missing the following information:

- For global zones and zone clusters, the link aggregation that you create in Step 2 requires two links.
- For a zone cluster's interconnect, you must create a virtual NIC that uses the link aggregation and the virtual LAN ID (VID) and you must specify the resulting virtual NIC.

Add the following command after the `dladm show-phys` command in Step 2:

```
# dladm create-vnic -l aggr2 -v 2853 zc_aggr2_vnic1
```

- For global zones and zone clusters, you only need to create a single interconnect by configuring either a VLAN for global zones or a virtual NIC for the zone cluster. A second aggregation, VLAN/VNIC, is not required because the single interconnect uses the underlying aggregation for redundancy.

geopg(8) Man Page

The `geopg(8)` man page is missing information about the extension properties for Oracle GoldenGate data replication for Disaster Recovery Framework. See [Appendix A, Oracle GoldenGate Replication Extension Properties](#), in *Oracle Solaris Cluster Data Replication Guide for Oracle GoldenGate* for information about these extension properties.