

Oracle® Solaris Cluster 4.1 Release Notes

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

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

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

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

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

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

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

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

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

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

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

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

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

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

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

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

Contents

| | |
|---|----|
| Preface | 7 |
| 1 Oracle Solaris Cluster 4.1 Release Notes | 13 |
| What's New in the Software | 13 |
| Support for Oracle Solaris 11.2 OS | 14 |
| New clsetup Wizards to Create a Zone Cluster | 15 |
| Support for solaris10 Brand Zone Clusters | 15 |
| Support for Exclusive-IP Zone Clusters | 15 |
| Support for Trusted Extensions With Zone Clusters | 15 |
| Resource Dependencies Can Be Defined on a Per-Node Basis | 16 |
| Support for Kernel Cage Dynamic Reconfiguration (DR) | 16 |
| Cluster Security Framework Is Enhanced | 16 |
| Support for Socket Direct Protocol Over the Cluster Interconnect | 16 |
| Faster Failure Detection and Response by Storage Monitors | 16 |
| ZFS Storage Pools | 17 |
| New clsetup Wizard to Configure the Oracle PeopleSoft Application Server Data Service | 17 |
| New clsetup Wizard to Configure the Oracle WebLogic Server Data Service | 17 |
| Support for MySQL and MySQL Cluster Data Services | 18 |
| New Data Service for PostgreSQL | 18 |
| New Data Service for Samba | 18 |
| New Data Service for SAP liveCache | 18 |
| New Data Service for SAP MaxDB | 18 |
| New Data Service for Siebel 8.2.2 | 18 |
| New Data Service for Sybase ASE | 19 |
| New Data Service for Oracle Traffic Director | 19 |
| New Data Service for Oracle TimesTen | 19 |
| New Manual for SAP NetWeaver Data Service | 19 |

| | |
|---|----|
| New Data Service for Oracle External Proxy | 19 |
| New Data Service for Oracle PeopleSoft Enterprise Process Scheduler | 19 |
| New Data Service for Oracle Web Tier | 20 |
| Support for Oracle E-Business 12.1.1 Data Service | 20 |
| Support for Sun ZFS Storage Appliance Data Replication With Geographic Edition | 20 |
| Support for EMC Symmetrix Remote Data Facility With Geographic Edition | 20 |
| Support for MySQL Replication With Geographic Edition | 20 |
| New Man Pages for the <code>ccradm</code> and <code>dcscfg</code> Advanced Maintenance Commands | 20 |
| Selected Support for Non-Global Zones | 21 |
| What's Not Included in the Oracle Solaris Cluster 4.1 Software | 21 |
| Restrictions | 22 |
| Solaris Volume Manager Disk Sets in a Zone Cluster | 22 |
| Commands Modified in This Release | 22 |
| Compatibility Issues | 22 |
| Logical Host Does not Fail Over with Public Net Fault (16979921) | 22 |
| Oracle ASM With Solaris Volume Manager Mirrored Logical Volumes | 23 |
| osysmond Core Dumps in S10 Brand Zone During <code>GI root.sh</code> and Starting of CRS (14456069) | 23 |
| Oracle Clusterware Fails to Create All SIDs for <code>ora.asm</code> Resource (12680224) | 23 |
| Oracle Solaris 11 SRU Installation Might Fail Due to Out-of-Date <code>pkg</code> Command | 24 |
| Unable to Install Just Patches Using <code>clzonecluster install-cluster</code> to the <code>solaris10</code> Branded Zone Cluster (7200532) | 24 |
| Adding Main Adapter to IPMP Group Removes DNS Configuration (7198718) | 24 |
| SAP JAVA Issue Affects HA for SAP NetWeaver Ability to Fail Over in Unplanned Outage (7191360) | 25 |
| Zone Does Not Boot if <code>pkg:/system/resource-mgmt/resource-cap</code> Is Not Installed and capped-memory Is Configured (7087700) | 25 |
| Active:Active ZFS Storage Appliance Clustered Configurations Are Not Supported With Geographic Edition (6770212) | 26 |
| Accessibility Information | 26 |
| Supported Products | 26 |
| Data Replication | 27 |
| Data Service | 27 |
| File Systems | 28 |
| Geographic Edition Software Requirements | 28 |
| Memory Requirements | 29 |
| Oracle Solaris Operating System | 29 |

| | |
|---|-----------|
| Oracle VM Server for SPARC | 29 |
| Volume Management Software | 29 |
| Product Localization | 30 |
| Known Issues and Bugs | 30 |
| Administration | 31 |
| Data Services | 36 |
| Developer Environment | 39 |
| Geographic Edition | 40 |
| Installation | 41 |
| Runtime | 45 |
| Upgrade | 46 |
| Software Updates | 47 |
| Patch Management Tools | 47 |
| My Oracle Support | 48 |
| Oracle Solaris Cluster 4.1 Documentation Set | 48 |
| Documentation Issues | 48 |
| Upgrade Guide | 48 |
| HA for Oracle Guide | 49 |
| HA for Oracle RAC Guide | 49 |
| HA for Oracle Solaris Zones Guide | 50 |
| Solaris Volume Manager | 50 |
| Geographic Edition Data Replication Guide for Oracle Solaris Availability Suite | 50 |
| Man Pages | 51 |
| A ORCL.otd(5) Man Page | 55 |
| ORCL.otd(5) Man Page | 55 |
| Description | 55 |
| Standard Properties | 56 |
| Extension Properties | 57 |
| Examples | 60 |
| Attributes | 62 |
| See Also | 62 |

Preface

This document lists product features, requirements, and restrictions for the Oracle Solaris Cluster 4.1 and Oracle Solaris Cluster Geographic Edition (Geographic Edition) 4.1 software on both SPARC based systems and x86 based systems. It also describes open defects and other known problems.

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

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

Bash is the default shell for Oracle Solaris 11. Machine names shown with the Bash shell prompt are displayed for clarity.

Using UNIX Commands

This document contains information about commands that are specific to installing and configuring Oracle Solaris Cluster data services. The document does *not* contain comprehensive information about basic UNIX commands and procedures, such as shutting down the system, booting the system, and configuring devices. Information about basic UNIX commands and procedures is available from the following sources:

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

Typographic Conventions

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

TABLE P-1 Typographic Conventions

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

Shell Prompts in Command Examples

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

TABLE P-2 Shell Prompts

| Shell | Prompt |
|--|----------------------------|
| Bash shell, Korn shell, and Bourne shell | \$ |
| Bash shell, Korn shell, and Bourne shell for superuser | # |
| C shell | <code>machine_name%</code> |
| C shell for superuser | <code>machine_name#</code> |

Related Documentation

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

| Topic | Documentation |
|--|---|
| Hardware installation and administration | <i>Oracle Solaris Cluster 4.1 Hardware Administration Manual</i> Individual hardware administration guides |
| Concepts | <i>Oracle Solaris Cluster Concepts Guide</i> |
| Software installation | <i>Oracle Solaris Cluster Software Installation Guide</i> |
| Data service installation and administration | <i>Oracle Solaris Cluster Data Services Planning and Administration Guide</i> and individual data service guides |
| Data service development | <i>Oracle Solaris Cluster Data Services Developer's Guide</i> |
| System administration | <i>Oracle Solaris Cluster System Administration Guide</i> <i>Oracle Solaris Cluster Quick Reference</i> |
| Software upgrade | <i>Oracle Solaris Cluster Upgrade Guide</i> |
| Error messages | <i>Oracle Solaris Cluster Error Messages Guide</i> |
| Command and function references | <i>Oracle Solaris Cluster Reference Manual</i> <i>Oracle Solaris Cluster Data Services Reference Manual</i> <i>Oracle Solaris Cluster Geographic Edition Reference Manual</i> <i>Oracle Solaris Cluster Quorum Server Reference Manual</i> |
| Compatible software | Oracle Solaris Cluster Compatibility Guide available at the Oracle Solaris Cluster Technical Resources page |

Access to Oracle Support

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

Getting Help

If you have problems installing or using Oracle Solaris Cluster, contact your service provider and provide the following information.

- Your name and email address (if available)
- Your company name, address, and phone number
- The model number and serial number of your systems
- The release number of the operating environment (for example, Oracle Solaris 11)
- The release number of Oracle Solaris Cluster (for example, Oracle Solaris Cluster 4.1)

Use the following commands to gather information about your system for your service provider.

| Command | Function |
|--|---|
| <code>prtconf -v</code> | Displays the size of the system memory and reports information about peripheral devices |
| <code>psrinfo -v</code> | Displays information about processors |
| <code>pkg list</code> | Reports which packages are installed |
| <code>prtdiag -v</code> | Displays system diagnostic information |
| <code>/usr/cluster/bin/clnode show-rev -v</code> | Displays Oracle Solaris Cluster release and package version information for each node |

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

Revision History

The following table lists the information that has been revised or added since the initial release of this documentation. The table also lists the revision date for these changes.

| Revision Date | New Information |
|---------------|--|
| November 2012 | Added support for the Oracle RAC data service in the global cluster. |

| Revision Date | New Information |
|---------------|---|
| February 2013 | <p>Added support for the Siebel data service.</p> <p>Added Oracle Solaris ZFS Volume Manager as supported with Sun ZFS Storage Appliance from Oracle data replication.</p> <p>Removed the erroneous requirement that Geographic Edition requires Oracle Solaris version 11.1.</p> <p>Added a command syntax correction for the Upgrade Guide.</p> |
| July 2013 | <p>Added support for new data services (PostgreSQL, SAP liveCache, SAP MaxDB, Siebel, Sybase ASE, Samba, and Oracle Traffic Director) to the What's New section.</p> <p>Updated the Compatibility section with information for Oracle ASM with Solaris Volume Manager Mirrored Logical Volumes.</p> <p>Updated the Known Issues section.</p> |

Oracle Solaris Cluster 4.1 Release Notes

This document provides the following information for the Oracle Solaris Cluster 4.1 software and the Oracle Solaris Cluster Geographic Edition 4.1 software, as of the time of publication.

- “What's New in the Software” on page 13
- “What's Not Included in the Oracle Solaris Cluster 4.1 Software” on page 21
- “Restrictions” on page 22
- “Commands Modified in This Release” on page 22
- “Compatibility Issues” on page 22
- “Accessibility Information” on page 26
- “Supported Products” on page 26
- “Product Localization” on page 30
- “Known Issues and Bugs” on page 30
- “Software Updates” on page 47
- “Oracle Solaris Cluster 4.1 Documentation Set” on page 48
- “Documentation Issues” on page 48

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

What's New in the Software

This section provides information related to new features, functionality, and products in the Oracle Solaris Cluster 4.1 and Oracle Solaris Cluster Geographic Edition (Geographic Edition) 4.1 software.

The Oracle Solaris Cluster 4.1 software provides the following new features:

- “Support for Oracle Solaris 11.2 OS” on page 14
- “New `clsetup` Wizards to Create a Zone Cluster” on page 15

- “Support for solaris10 Brand Zone Clusters” on page 15
- “Support for Exclusive-IP Zone Clusters” on page 15
- “Support for Trusted Extensions With Zone Clusters” on page 15
- “Resource Dependencies Can Be Defined on a Per-Node Basis” on page 16
- “Support for Kernel Cage Dynamic Reconfiguration (DR)” on page 16
- “Cluster Security Framework Is Enhanced” on page 16
- “Support for Socket Direct Protocol Over the Cluster Interconnect” on page 16
- “Faster Failure Detection and Response by Storage Monitors” on page 16
- “ZFS Storage Pools” on page 17
- “New clsetup Wizard to Configure the Oracle PeopleSoft Application Server Data Service” on page 17
- “New clsetup Wizard to Configure the Oracle WebLogic Server Data Service” on page 17
- “Support for MySQL and MySQL Cluster Data Services” on page 18
- “New Data Service for PostgreSQL” on page 18
- “New Data Service for Samba” on page 18
- “New Data Service for SAP liveCache” on page 18
- “New Data Service for SAP MaxDB” on page 18
- “New Data Service for Siebel 8.2.2” on page 18
- “New Data Service for Sybase ASE” on page 19
- “New Data Service for Oracle Traffic Director” on page 19
- “New Data Service for Oracle TimesTen” on page 19
- “New Manual for SAP NetWeaver Data Service” on page 19
- “New Data Service for Oracle External Proxy” on page 19
- “New Data Service for Oracle PeopleSoft Enterprise Process Scheduler” on page 19
- “New Data Service for Oracle Web Tier” on page 20
- “Support for Oracle E-Business 12.1.1 Data Service” on page 20
- “Support for Sun ZFS Storage Appliance Data Replication With Geographic Edition” on page 20
- “Support for EMC Symmetrix Remote Data Facility With Geographic Edition” on page 20
- “Support for MySQL Replication With Geographic Edition” on page 20
- “New Man Pages for the ccradm and dcs_config Advanced Maintenance Commands” on page 20
- “Selected Support for Non-Global Zones” on page 21

Support for Oracle Solaris 11.2 OS

Support is added for Oracle Solaris 11.2 software. The cluster must run or be upgraded to a minimum of Oracle Solaris Cluster 4.1 SRU 8.

New clsetup Wizards to Create a Zone Cluster

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

Support for solaris10 Brand Zone Clusters

Support is added for `solaris10` brand zone clusters. At initial release of this feature, the following Oracle Solaris Cluster 3.3 data services are supported in `solaris10` brand zone clusters on an Oracle Solaris Cluster 4.1 configuration:

- Oracle RAC for the following releases:
 - 11g release 2, on SPARC (without UDLM) or on x86
 - 10g release 2, on x86 only
- HA for Oracle PeopleSoft, on SPARC only
- HA for Oracle WebLogic Server, on SPARC or x86

Oracle Solaris Cluster Geographic Edition software is **not** initially supported in `solaris10` brand zone clusters.

To learn whether additional Oracle Solaris Cluster 3.3 data services or the Geographic Edition feature become qualified in an Oracle Solaris Cluster 4.1 `solaris10` brand zone cluster, consult the [Oracle Solaris Cluster 4 Compatibility Guide](#).

For instructions to create a `solaris10` brand zone cluster, see [“How to Create a Zone Cluster”](#) in *Oracle Solaris Cluster Software Installation Guide*.

Support for Exclusive-IP Zone Clusters

Exclusive-IP zone clusters work with `solaris` and `solaris10` brand zones. For installation instructions of `solaris` brand exclusive-IP zone clusters, see [“How to Create a Zone Cluster”](#) in *Oracle Solaris Cluster Software Installation Guide*. For `solaris10` brand exclusive-IP zone clusters, see the `readme` file of the SRU3.

Support for Trusted Extensions With Zone Clusters

You can now configure the Trusted Extensions feature of Oracle Solaris software for use with Zone Cluster. For details, see [“How to Create a Zone Cluster”](#) in *Oracle Solaris Cluster Software Installation Guide*.

Resource Dependencies Can Be Defined on a Per-Node Basis

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

Support for Kernel Cage Dynamic Reconfiguration (DR)

Kernel cage memory DR (the removal of system boards containing system memory) is now supported. Support is currently limited to Oracle's SPARC Enterprise M8000 and M9000 series servers. For details, see [“Kernel Cage DR Recovery”](#) in *Oracle Solaris Cluster 4.1 Hardware Administration Manual*.

Cluster Security Framework Is Enhanced

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

Support for Socket Direct Protocol Over the Cluster Interconnect

Socket Direct Protocol is now supported over the Oracle Solaris Cluster interconnect. For details, see [“Requirements for Socket Direct Protocol Over an Oracle Solaris Cluster Interconnect”](#) in *Oracle Solaris Cluster 4.1 Hardware Administration Manual*.

Faster Failure Detection and Response by Storage Monitors

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

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

See the following man pages for details:

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

ZFS Storage Pools

Support is added for ZFS storage pools on replicated storage that is managed by Oracle Solaris Cluster software. The following data-replication software is supported for this feature:

- EMC Symmetrix Remote Data Facility (SRDF)
- Sun ZFS Storage Appliance

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

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

New clsetup Wizard to Configure the Oracle WebLogic Server Data Service

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

Support for MySQL and MySQL Cluster Data Services

Data services for MySQL and MySQL Cluster are now supported. For details, see the following:

- [Oracle Solaris Cluster Data Service for MySQL Guide](#)
- [Oracle Solaris Cluster Data Service for MySQL Cluster Guide](#)

New Data Service for PostgreSQL

A new data service for PostgreSQL software is now available with Oracle Solaris Cluster 4.1 SRU3. For details, see [Oracle Solaris Cluster Data Service for PostgreSQL Guide](#).

New Data Service for Samba

A new data service for Samba 3.6.6 and 3.6.12 software is now available with Oracle Solaris Cluster 4.1 SRU3. For details, see *Deployment Example: Configuring HA Samba on Oracle Solaris 11.1* (Article 1509776) posted on My Oracle Support at <http://support.oracle.com>. You must have an Oracle support contract to access the site.

New Data Service for SAP liveCache

A new data service for SAP liveCache software is now available with Oracle Solaris Cluster 4.1 SRU3. For details, see [Oracle Solaris Cluster Data Service for SAP liveCache Guide](#).

New Data Service for SAP MaxDB

A new data service for SAP MaxDB software is now available with Oracle Solaris Cluster 4.1 SRU3. For details, see [Oracle Solaris Cluster Data Service for SAP MaxDB Guide](#).

New Data Service for Siebel 8.2.2

A new data service for Siebel 8.2.2 software is now available with Oracle Solaris Cluster 4.1 SRU3. For details, see *Configuring the Oracle Solaris Cluster Data Service for Siebel* (Article 1509776) posted on My Oracle Support at <http://support.oracle.com>. You must have an Oracle support contract to access the site.

New Data Service for Sybase ASE

A new data service for Sybase ASE 15.7 is now available with Oracle Solaris Cluster 4.1 SRU3. For details, see the *Deployment Example: Configuring the Oracle Solaris Cluster HA for Sybase Adaptive Server Enterprise (ASE) 15.7 Agent* (Article 1509776) posted on My Oracle Support at <http://support.oracle.com>. You must have an Oracle support contract to access the site.

New Data Service for Oracle Traffic Director

A new data service for Oracle Traffic Director software is now available with Oracle Solaris Cluster 4.1 SRU3. For details, see the `ORCL.otd(5)` man page located in [Appendix A](#) of the Release Notes. See also Article 1509776 posted on My Oracle Support at <http://support.oracle.com>. You must have an Oracle support contract to access the site.

Support is currently limited to engineered systems. For more information, see [Oracle Solaris Cluster 4 Compatibility Guide](#) (<http://www.oracle.com/technetwork/server-storage/solaris-cluster/overview/solariscluster4-compatibilityguide-1429037.pdf>).

New Data Service for Oracle TimesTen

A new data service for Oracle TimesTen 11.2.2 software is now available with Oracle Solaris Cluster 4.1 SRU3. For details, see *Oracle Solaris Cluster Data Service for Oracle TimesTen Guide*.

New Manual for SAP NetWeaver Data Service

A new manual for the SAP NetWeaver 7.3 data service is now available with Oracle Solaris Cluster 4.1 SRU3. For details, see *Oracle Solaris Cluster Data Service for SAP NetWeaver Guide*.

New Data Service for Oracle External Proxy

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

New Data Service for Oracle PeopleSoft Enterprise Process Scheduler

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

New Data Service for Oracle Web Tier

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

Support for Oracle E-Business 12.1.1 Data Service

The data service for Oracle E-Business Suite 12.1.1 is now supported. For details, see [Oracle Solaris Cluster Data Service for Oracle E-Business Suite Guide](#).

Support for Sun ZFS Storage Appliance Data Replication With Geographic Edition

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

Support for EMC Symmetrix Remote Data Facility With Geographic Edition

EMC Symmetrix Remote Data Facility (SRDF) data replication for Geographic Edition software is now supported. For details, see [Oracle Solaris Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility](#).

Support for MySQL Replication With Geographic Edition

MySQL data replication for Geographic Edition software is now supported. For details, see Appendix F, “Deployment Example: Replicating Data With MySQL,” in [Oracle Solaris Cluster Geographic Edition System Administration Guide](#).

New Man Pages for the ccradm and dcs_config Advanced Maintenance Commands

The following man pages for advanced cluster administration tools are added in this release:

- [ccradm\(1M\)](#) – This man page describes the ccradm tool, which is an advanced maintenance command for administering the Cluster Configuration Repository (CCR).

- **dcscfg(1M)** – This man page describes the `dcscfg` tool, which is an advanced maintenance command for administering the Device Configuration System (DCS). This command was previously only available as part of a diagnostic toolkit but is now added to Oracle Solaris Cluster software in the 4.1 release.

Selected Support for Non-Global Zones

The Oracle Solaris Cluster 4.1 release runs only in a global zone and in a zone cluster. A zone cluster is now configured with the `solaris` brand non-global zone, which is comparable to the native brand zone in Oracle Solaris 10.

The `solaris` and `solaris10` brands of non-global zones are supported for configuration with the HA for Oracle Solaris Zones data service.

In this release, the global-cluster nonvoting node model, to configure non-global zones in a resource-group node list, is not supported.

What's Not Included in the Oracle Solaris Cluster 4.1 Software

The following features are included in the Oracle Solaris Cluster 3.3 version but are not included in the Oracle Solaris Cluster 4.1 release. These features might become available at a later time:

- Support for Veritas File System (VxFS) and Veritas Volume Manager (VxVM)
- Support for the VxVM cluster feature for Oracle RAC in addition to VxVM with Oracle Solaris Cluster
- Support for Oracle Automatic Storage Management Cluster File System (Oracle ACFS)
- GUI and GUI wizards
- Support for Sun Management Center
- Support for Sun QFS from Oracle
- Support for non-global zones as resource-group node-list targets
- Support for Oracle Solaris IP Security Architecture (IPsec) on the private interconnect
- The `scsnapshot` tool
- The `cconsole` utility (the Oracle Solaris `pconsole` utility can be used instead)
- Support for Hitachi True Copy and Hitachi Universal Replicator storage-based replication

The following HA data services are not available at the time of publication but might become available at a later time:

- Afsa IMPAX

- Informix
- Kerberos
- Oracle Application Server
- Oracle Business Intelligence Enterprise Edition
- Oracle iPlanet Web Proxy Server
- SAP Web Application Server
- SWIFTAlliance Access and Gateway
- WebSphere Message Broker
- WebSphere Message Queue

Restrictions

The following restriction is in effect as of the time of publication.

Contact your Oracle representative to learn whether a feature or product becomes qualified with Oracle Solaris Cluster 4.1 software at a later date.

Solaris Volume Manager Disk Sets in a Zone Cluster

Oracle Solaris Bug 15764835 prevents the use of Solaris Volume Manager multi-owner disk sets on zone-cluster nodes. Attempting to use multi-owner disk sets in a zone cluster might cause the underlying global-cluster nodes to panic.

Commands Modified in This Release

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

Compatibility Issues

This section contains information about Oracle Solaris Cluster compatibility issues with other products, as of initial release. Contact your Oracle support representative to see whether a fix becomes available.

Logical Host Does not Fail Over with Public Net Fault (16979921)

Problem Summary: IPMP groups in exclusive-IP zone clusters fail to recognize link failures that cause dependent logical hostname resources to remain online, even when the base network interface link is broken.

Workaround: Enable transitive probing for the IPMP network service or create probe-based IPMP groups in exclusive-IP zone clusters.

Oracle ASM With Solaris Volume Manager Mirrored Logical Volumes

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

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

osysmond Core Dumps in S10 Brand Zone During `Grid root.sh` and Starting of CRS (14456069)

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

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

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

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

```
ORACLE_SID (+ASM2) does not match the Oracle ASM configuration ORACLE_SID () within CRS
```

```
ERROR: Oracle ASM is either not installed or the installation is invalid!
```

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

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


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

Oracle Solaris 11 SRU Installation Might Fail Due to Out-of-Date pkg Command

Problem Summary: When you install an Oracle Solaris 11 SRU to your cluster prior to the upgrade to Oracle Solaris 11.1, you might receive an error message similar to the following:

```
WARNING: pkg(5) appears to be out of date, and should be updated before  
running update. Please update pkg(5) by executing 'pkg install  
pkg:/package/pkg' as a privileged user and then retry the update.
```

Workaround: Follow the instructions in the error message.

Unable to Install Just Patches Using clzonecluster install-cluster to the solaris10 Branded Zone Cluster (7200532)

Problem Summary: The clzonecluster install-cluster command might fail to install a patch on a solaris10 brand zone if Oracle Solaris Cluster patch 145333-15 (SPARC) or 145334-15 (x86) is installed in the zone. For example:

```
# clzonecluster install-cluster -p patchdir=/var/tmp/patchdir,patchlistfile=plist S10ZC  
Installing the patches ...  
clzc: (C287410) Failed to execute command on node "zcnodel":  
scpatchadm: Logging reports to "/var/cluster/logs/install/scpatchadm.log.123"  
  
scpatchadm.log.123 would show the message:  
scpatchadm: Failed to install the following patches:  
123456-01  
clzc: (C287410) Failed to execute command on node "zcnodel"
```

Workaround: Log in to the zone and install the patch by using the patchadd command.

Contact your Oracle support representative to learn whether an Oracle Solaris Cluster 3.3 patch becomes available.

Adding Main Adapter to IPMP Group Removes DNS Configuration (7198718)

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


```
# ipadm delete-ip adapter
# ipadm create-ip adapter
# ipadm create-imp -i adapter sc_imp0
# ipadm create-addr -T static -a local=hostname/24 sc_imp0/v4
```

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

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

SAP JAVA Issue Affects HA for SAP NetWeaver Ability to Fail Over in Unplanned Outage (7191360)

The SAP JAVA stack has a severe problem that affects the failover of dialogue instances in an HA for SAP NetWeaver configuration. On an unplanned node outage, like a panic or power outage, the SAP message server does not accept the connection of a dialogue instance on a different node until a timeout is over. This leads to the following behavior:

Once a node that is hosting a failover dialogue instance panics or experiences an outage, the dialogue instance does not start on the target node on the first try. The dialogue instance will do one of the following:

- Come online after one or more retries.
- Fail back to the original node if that node comes back up early enough.

This behavior occurs only on unplanned outages. Any orderly shutdown of a node does not experience this problem. Also, ABAP or dual-stack configurations are not affected.

Zone Does Not Boot if pkg:/system/resource-mgmt/resource-cap Is Not Installed and capped-memory Is Configured (7087700)

Problem Summary: If the package `pkg:/system/resource-mgmt/resource-cap` is not installed and a zone is configured with `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 `pkg:/system/resource-mgmt/resource-cap` into the global zone. Once the `resource-cap` package is installed, the zone can boot.

Active:Active ZFS Storage Appliance Clustered Configurations Are Not Supported With Geographic Edition (6770212)

At initial release of Oracle Solaris Cluster 4.1 software, an active:active remote replication in a clustered configuration, where both heads are replicating data, is not supported by Sun ZFS Storage Appliance. Contact your Oracle support representative to learn whether a patch or workaround is available.

However, active-passive configurations are currently supported in a clustered configuration.

Accessibility Information

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

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

Supported Products

This section describes the software and memory requirements for Oracle Solaris Cluster 4.1 and Oracle Solaris Cluster Geographic Edition 4.1 software, as of the time of publication. For the latest information about supported products and product versions for this release, see the [Oracle Solaris Cluster 4 Compatibility Guide \(http://www.oracle.com/technetwork/server-storage/solaris-cluster/overview/solariscluster4-compatibilityguide-1429037.pdf\)](http://www.oracle.com/technetwork/server-storage/solaris-cluster/overview/solariscluster4-compatibilityguide-1429037.pdf).

- “Data Replication” on page 27
- “Data Service” on page 27
- “File Systems” on page 28
- “Geographic Edition Software Requirements” on page 28
- “Memory Requirements” on page 29
- “Oracle Solaris Operating System” on page 29
- “Oracle VM Server for SPARC” on page 29
- “Volume Management Software” on page 29

Data Replication

The following data replication products are supported in this release:

- **Availability Suite feature of Oracle Solaris**
- **EMC Symmetrix Remote Data Facility (SRDF)**
- **Oracle Data Guard** – Support is limited to configurations that use Oracle Database 11g release 2 software.
- **Oracle Solaris Cluster Geographic Edition script-based plug-ins.**
- **Sun ZFS Storage Appliance** – When used with Geographic Edition software, support requires a minimum of Sun ZFS 7000 Storage Appliance 2011.1.5 software.

The following table describes the volume managers that are supported in this release for each data replication software product.

| Data Replication Software | Supported Volume Managers |
|--|--|
| Availability Suite feature of Oracle Solaris | Solaris Volume Manager |
| EMC Symmetrix Remote Data Facility | Solaris Volume Manager |
| Oracle Data Guard ¹ | Oracle Automatic Storage Management |
| | Oracle Solaris ZFS Volume Manager |
| | Solaris Volume Manager for Sun Cluster |
| Sun ZFS Storage Appliance from Oracle | Solaris Volume Manager |
| | Oracle Solaris ZFS Volume Manager |

¹ For information about additional supported storage management options, see “Storage Management Requirements” in *Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide*.

Data Service

At the time of publication, data services for the following applications are supported with Oracle Solaris Cluster 4.1:

- Apache
- Apache Tomcat
- DHCP
- DNS
- MySQL
- MySQL Cluster
- NFS
- Oracle eBusiness Suite
- Oracle External Proxy

- Oracle GlassFish Server Message Queue
- Oracle Database single-instance on version 11g release 2 only
- Oracle iPlanet Web Server
- Oracle PeopleSoft Enterprise
- Oracle Real Application Clusters on version 11g release 2 only
- Oracle Solaris Zones
- Oracle Traffic Director
- Oracle VM Server for SPARC
- Oracle Web Tier
- Oracle WebLogic Server
- PostgreSQL
- Samba
- SAP liveCache
- SAP MaxDB
- SAP NetWeaver
- Siebel
- Sybase ASE

File Systems

The following table describes the file systems that work with the Oracle Solaris Cluster 4.1 software on SPARC and x86 platforms.

| File System | Additional Information |
|---|---|
| Oracle Solaris UFS | |
| Oracle Solaris ZFS | |
| Network File System (NFS) | |
| Oracle Solaris Cluster proxy file system (PxFS) | Only a UFS file system can be configured as a cluster file system |

Geographic Edition Software Requirements

The Oracle Solaris Cluster Geographic Edition (Geographic Edition) 4.1 software runs only on either Oracle Solaris Cluster software version 4.1 or 4.0.

A cluster running Geographic Edition 4.1 software can be in a partnership only with another cluster that runs either Geographic Edition software version 4.1 or 4.0.

Memory Requirements

The Oracle Solaris Cluster 4.1 software requires the following memory requirements for every cluster node:

- Minimum of 1.5 Gbytes of physical RAM (2 Gbytes typical)
- Minimum of 6 Gbytes of available hard drive space

Actual physical memory and hard drive requirements are determined by the applications that are installed. Consult the application's documentation or contact the application vendor to calculate additional memory and hard drive requirements.

Oracle Solaris Operating System

The Oracle Solaris Cluster 4.1 and Quorum Server software require the Oracle Solaris 11.1 OS. The Oracle Solaris 11.0 release is not supported for Oracle Solaris Cluster 4.1 software.

In addition, to upgrade directly to Oracle Solaris 11.1 software, before upgrade the system must run a minimum of Oracle Solaris 11 with SRU 10. Oracle Solaris software does not support direct upgrade to Oracle Solaris 11.1 from less than Oracle Solaris 11 with SRU 10.

Oracle VM Server for SPARC

This Oracle Solaris Cluster release works with Oracle VM Server for SPARC 2.1, 2.2, and 3.0. This software was formerly called Sun Logical Domains (LDoms).

The following describe the supported and restricted uses of the SR-IOV feature of Oracle Solaris 11.1 for use with Oracle VM Server for SPARC 2.2:

- An SR-IOV device is not supported for *live migration* of the HA for Oracle VM Server for SPARC data service. This restriction exists as of initial release of the Oracle Solaris Cluster 4.1 software. Contact your Oracle support representative to learn whether support for the SR-IOV feature becomes available.
- An SR-IOV device is supported for *normal* migration of the HA for Oracle VM Service for SPARC data service. This support requires that the SR-IOV device path is identical on all potential primaries of the data service's resource group.
- An SR-IOV device is supported with an Oracle VM Server for SPARC logical domain that is configured to run as a cluster node.

Volume Management Software

The following table describes the volume managers that work with Oracle Solaris 11.1 on the SPARC and x86 platforms. Veritas Volume Manager (VxVM) is not supported for the Oracle Solaris Cluster and Oracle Solaris Cluster Geographic Edition 4.1 software.

| Volume Manager | Cluster Feature |
|-----------------------------------|--|
| Solaris Volume Manager | Solaris Volume Manager for Sun Cluster |
| Oracle Solaris ZFS Volume Manager | not applicable |

Product Localization

The following table describes localization for certain components of Oracle Solaris Cluster 4.1 software:

| Component | Localization |
|-----------------------|--------------------------------------|
| Software Command Line | Japanese, Korean, Simplified Chinese |
| Man Pages | Japanese, Simplified Chinese |

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

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

Known Issues and Bugs

The following known issues and bugs affect the operation of the Oracle Solaris Cluster and Oracle Solaris Cluster Geographic Edition 4.1 software, as of the time of release. Bugs and issues are grouped into the following categories:

- [“Administration” on page 31](#)
- [“Data Services” on page 36](#)
- [“Developer Environment” on page 39](#)
- [“Geographic Edition” on page 40](#)
- [“Installation” on page 41](#)
- [“Runtime” on page 45](#)
- [“Upgrade” on page 46](#)

Contact your Oracle support representative to see whether a fix becomes available.

Administration

A `clzc` reboot Command Causes the solaris10 Brand Exclusive-IP Zone Cluster to Panic the Global Zone Nodes (16941521)

Problem Summary: A reboot or halt of a solaris10 branded exclusive-IP zone cluster node can cause the global zone nodes to panic. This occurs when the zone cluster nodes use the base network as the primary (public) network interface and there are VNICs on that base network interface that are configured for other zone cluster nodes in that cluster.

Workaround: Create and use VNICs as primary network interfaces for exclusive-IP zone clusters.

The `/usr/sbin/shutdown` Command in a Zone of an Exclusive-IP Zone Cluster Can Result in a Halt of Other Running Zones of the Zone Cluster (16963753)

Problem Summary: If you use the `/usr/sbin/shutdown` command in a zone of an exclusive-IP zone cluster to halt or reboot the zone, any other zones of the zone cluster that are alive and running can be halted by cluster software.

Workaround: Do not use the `/usr/sbin/shutdown` command inside a zone of an exclusive-IP zone cluster to halt or reboot the zone. Instead, use the `/usr/cluster/bin/clzonecluster` command in the global zone to halt or reboot a zone of an exclusive-IP zone cluster. The `/usr/cluster/bin/clzonecluster` command is the correct way to halt or reboot a zone of any type of zone cluster. If you see this problem, use the `/usr/cluster/bin/clzonecluster` command to boot any such zones that were halted by cluster software.

The `svc_private_network:default` SMF Service Goes Into Maintenance in a solaris10 Brand Exclusive-IP Zone Cluster (16716992)

Problem Summary: When you perform system identification in a zone of a solaris10 brand exclusive-IP zone cluster, the `svc_private_network:default` SMF service goes into maintenance in that zone. On subsequent reboots of the zone, the problem does not occur.

Workaround: After you perform system identification configuration in a zone of a solaris10 brand exclusive-IP zone cluster, reboot that zone.

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

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

Workaround: There is no known workaround.

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

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

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

The Data Service Configuration Wizards Do Not Support Storage Resources and Resource Groups for Scalable HASToragePlus (7202824)

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

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

Workaround: Configure data services that use scalable HASToragePlus resources and resource groups in the following way:

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

Removing a Node From an Exclusive-IP Zone Cluster Panics Cluster Nodes (7199744)

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 panic. 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.

Nonexisting privnet Stops Zone Clusters From Booting Despite Good privnet (7199431)

Problem Summary: If invalid or nonexisting network links are specified as the `privnetresources` in an exclusive-IP zone cluster configuration (`ip-type=exclusive`), the zone-cluster node fails to join the zone cluster despite presence of valid `privnet` resources.

Workaround: Remove the invalid `privnet` resource from the zone cluster configuration, then reboot the zone-cluster node.

```
# clzonecluster reboot -n nodename zone-cluster
```

Alternatively, create the missing network link that corresponds to the invalid `privnet` resource, then reboot the zone. See the [dladm\(1M\)](#) man page for more information.

The clzonecluster Command Fails to Verify That defrouter Cannot Be Specified Without allowed-addr, CCR Has Failed Configuration (7199135)

Problem Summary: In an exclusive-IP zone cluster, if you configure a `net` resource in the node scope with the `defrouter` property specified and the `allowed-address` property unspecified, the Oracle Solaris software errors out. Oracle Solaris software requires that, for an exclusive-IP zone cluster, you must always specify `allowed-address` property if you specify the `defrouter` property. If you do not, the Oracle Solaris software reports the proper error message, but the cluster would have already populated the CCR with the zone-cluster information. This action leaves the zone cluster in the Unknown state.

Workaround: Specify the `allowed-address` property for the zone cluster.

clzonecluster boot, reboot, and halt Subcommands Fail if Any One of the Cluster Nodes Is Not in the Cluster (7193998)

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

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



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

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

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

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

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

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

Cluster File System Does Not Support Extended Attributes (7167470)

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

- The extended attribute operations on a regular file will fail with a `ENOENT` error.
- The extended attribute operations on a directory will result as normal operations on the directory itself.

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

Workaround: Mounted a cluster file system with the `noxattrmount` option.

Using `chmod` to Set `setuid` Permission Returns Error in a Non-Global Zone on PxFS Secondary Server (7020380)

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

For example:

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

Workaround: Do one of the following:

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

Cannot Create a Resource From a Configuration File With Non-Tunable Extension Properties (6971632)

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

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

Disabling Device Fencing While Cluster Is Under Load Results in Reservation Conflict (6908466)

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

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

EMC SRDF Rejects Switchover When Replicated Device-Group Status Will Cause Switchover and Switchback to Fail (6798901)

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

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

Removing Nodes From the Cluster Configuration Can Result in Node Panics (6735924)

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

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

More Validation Checks Needed When Combining DIDs (6605101)

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

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

Data Services

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

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

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

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

Failure to Update Properties of SUNW.ScalMountPoint Resource Configured with NAS for Zone Cluster (7203506)

Problem Summary: The update of any properties in a `SUNW.ScalMountPoint` resource that is configured with a NAS file system for a zone cluster can fail with an error message similar to the following:

```
clrs:  hostname:zone-cluster : Bad address
```

Workaround: Use the `clresource` command to delete the resource and then recreate resource with all required properties.

Global File System Configured in Zone Cluster's Scalable HASToragePlus Resource Is Not Accessible (7197623)

Problem Summary: Consider a cluster file system with the following entry in the global cluster's `/etc/vfstab` file, with a `mount-at-boot` value of `no`:


```
# cat /etc/vfstab
/dev/md/datadg/dsk/d0    /dev/md/datadg/rdisk/d0 /global/fs-data ufs    5    no    logging,global
```

When an HAStoragePlus resource is created in a zone cluster's scalable resource group and the above cluster file system has the mount-at-boot value set to no, the cluster file system data might not be visible through the zone-cluster node mount point.

Workaround: Perform the following steps to avoid the problem:

1. From one global-cluster node, take offline the zone cluster's scalable resource group that contains HAStoragePlus.


```
# clresourcegroup offline -Z zonecluster scalable-resource-group
```
2. In the /etc/vfstab file on each global-cluster node, change the mount-at-boot value of the cluster file system entry to yes.

```
/dev/md/datadg/dsk/d0    /dev/md/datadg/rdisk/d0 /global/fs-data ufs    5    yes   logging,global
```

3. From one global-cluster node, bring online the zone cluster's scalable resource group that contains HAStoragePlus.

```
# clresourcegroup online -Z zonecluster scalable-resource-group
```

RAC Wizard Failing With "ERROR: Oracle ASM is either not installed or the installation is invalid!" (7196184)

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

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

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

clsetup Wizard Fails While Configuring WebLogic Server Domain in the Zones/Zone Cluster With WebLogic Server Installed in the NFS (7196102)

Problem Summary: The configuration of the HA-WebLogic Server resource using the clsetup wizard inside a zone/zone cluster would fail if the WebLogic Server is installed on an NFS mount point.

This issue won't occur with the NFS storage on global cluster, and if storage other than NFS is used.

Condition for this issue to occur : Mount the NFS storage with WebLogic Server installed inside the zones and configure the WebLogic Server using the clsetup wizard.

Error Message : ERROR: The specified path is not a valid WebLogic Server domain location. Similar message will be displayed for Home Location, Start Script and Environment file

Finally it fails in Administration/Managed/RPS server discovery.

Not able to find the WebLogic Administration Server Instance.
Make sure the provided WebLogic Domain Location (<DOMAIN_LOCATION_PROVIDED>) is the valid one.

No Reverse Proxy Server Instances found. You can't proceed further.

No Managed Server instances found. You can't proceed further.

Workaround: Configure the WebLogic Server resource manually.

With a Large Number of Non-Network-Aware GDS Resources, Some Fail to Restart and Remain Offline (7189659)

Problem Summary: This problem affects Generic Data Service (GDS) resources that meet all of the following conditions:

- No custom probe script is configured
- The `network_aware` property is set to FALSE.
- The `Retry_count` property is set to -1.

If the resources continue to fail to start, GDS will continue to restart it, forever. There is an issue where the error Restart operation failed: cluster is reconfiguring is produced. This results in the GDS resource not being automatically restarted.

Workaround: Manually disable and then re-enable the affected GDS resources.

SUNW.Proxy_SMF_failover sc_delegated_restarter File Descriptor Leak (7189211)

Problem Summary: Every time the SMF proxy resource `SUNW.Proxy_SMF_failover` is disabled or enabled, the file descriptor count increases by one. Repeated switches can grow the file descriptors to 256 and reach the limit at which point the resource cannot be switched online anymore.

Workaround: Disable and re-enable the `sc_restarter` SMF service.

```
# svcadm disable sc_restarter
# svcadm enable sc_restarter
```


When set `Debug_level=1`, `pas-rg` Fails Over to Node 2 And Cannot Start on Node 1 Anymore (7184102)

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

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

Scalable Applications Are Not Isolated Between Zone Clusters (6911363)

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

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

Running `clnas add` or `clnas remove` Command on Multiple Nodes at the Same Time Could Cause Problem (6791618)

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

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

Developer Environment

`clresource show -p` Command Returns Wrong Information (7200960)

Problem Summary: In a `solaris10` brand non-global zone, the `clresource show -p` *property* command returns the wrong information.

Workaround: This bug is caused by pre-Oracle Solaris Cluster 4.1 binaries in the `solaris10` brand zone. Run the following command from the global zone to get the correct information about local non-global zone resources:

```
# clresource show -p property -Z zone-name
```


Geographic Edition

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

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

Workaround: Reboot the node.

DR State Stays Reporting unknown on One Partner (7189050)

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

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

Takeover to the Secondary Is Failing Because fs umount Failed On the Primary (7182720)

Problem Summary: Takeover of a protection group fails if umount of the file system fails on the primary site.

Workaround: Perform the following steps:

1. Issue **fuser -cu file-system**.
2. Check for non-application process IDs, like **cd**, on the primary site.
3. Terminate such processes before you perform a takeover operation.

ZFS Storage Appliance Protection Group Creation And Validation Fail if Project Replication Is Stopped by Using the BUI (7176292)

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

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

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

Multiple Notification Emails Sent From Global Cluster When Zone Clusters Are in Use (7098290)

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

Workaround: This is a side effect of the cluster event handling. It is harmless, and the duplicates should be ignored.

Installation

Unable to Install Data Service Agents on Existing 3.3 5/11 solaris10 Brand Zone Without Specifying Patch Options (7197399)

Problem Summary: When installing agents in a solaris10 brand non-global zone from an Oracle Solaris Cluster 3.3 or 3.3 5/11 DVD, the `clzonecluster install-cluster` command fails if you do not specify the patches that support solaris10 branded zones.

Workaround: Perform the following steps to install agents from an Oracle Solaris Cluster 3.3 or 3.3 5/11 DVD to a solaris10 brand zone:

1. Reboot the zone cluster into offline mode.

```
# clzonecluster reboot -o zonecluster
```

2. Run the `clzonecluster install-cluster` command, specifying the information for the core patch that supports solaris10 branded zones.

```
# clzonecluster install-cluster -d dvd -p patchdir=patchdir[,patchlistfile=patchlistfile] \
-n node[,...] zonecluster
```

3. After installation is complete, reboot the zone cluster to bring it online.

```
# clzonecluster reboot zonecluster
```

clzonecluster Does Not Report Errors When install Is Used Instead of install-cluster for solaris10 Branded Zones (7190439)

Problem Summary: When the `clzonecluster install` command is used to install from an Oracle Solaris Cluster release DVD, it does not print any messages but nothing is installed onto the nodes.

Workaround: To install the Oracle Solaris Cluster release in a solaris10 branded zone, do not use the `clzonecluster install` command, which is used to install the Oracle Solaris 10 image. Instead, use the `clzonecluster install-cluster` command.

ASM Instance Proxy Resource Creation Errored When a Hostname Has Uppercase Letters (7190067)

Problem Summary: The use of uppercase letters in the cluster node hostname causes the creation of ASM instance proxy resources to fail.

Workaround: Use only lowercase letters for the cluster-node hostnames when installing Oracle Solaris Cluster software.

Wizard Won't Discover the ASM SID (7190064)

Problem Summary: When using the `clsetup` utility to configure the HA for Oracle or HA for Oracle RAC database, the Oracle ASM System Identifier screen is not able to discover or configure the Oracle ASM SID when a cluster node hostname is configured with uppercase letters.

Workaround: Use only lowercase letters for the cluster-node hostnames when installing Oracle Solaris Cluster software.

RAC Proxy Resource Creation Fails When the Cluster Node's Hostname Has Uppercase Letters (7189565)

Problem Summary: The use of uppercase letters in the cluster node hostname causes the creation of RAC database proxy resources to fail.

Workaround: Use only lowercase letters for the cluster-node hostnames when you install Oracle Solaris Cluster software.

Hard to Get Data Service Names for solaris10 Brand Zone Noninteractive Data Service Installation (7184714)

Problem Summary: It is hard to know what is the agent names to specify when using the `clzonecluster install-cluster` command to install agents with the `-s` option.

Workaround: When using the `clzonecluster install-cluster -d dvd -s {all | software-component[,...]} options zone-cluster` command to create a solaris10 brand 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)

cacao Cannot Communicate on Machines Running Trusted Extensions (7183625)

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

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

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

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

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

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

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

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

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

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

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

```
phys-schost-2# cd /etc/cacao/instances/default/
phys-schost-2# tar xf /tmp/SECURITY.tar
```

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

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

```
phys-schost-1# rm /tmp/SECURITY.tar
phys-schost-2# rm /tmp/SECURITY.tar
```

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

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

The Command `clnode remove -F nodename` Fails to Remove the Node *nodename* From Solaris Volume Manager Device Groups (6471834)

Problem Summary: When a node is removed from the cluster by using the command `clnode remove -F nodename`, a stale entry for the removed node might remain in Solaris Volume Manager device groups.

Workaround: Remove the node from the Solaris Volume Manager device group by using the `metaset` command before you run the `clnode remove -F nodename` command.

If you ran the `clnode remove -F nodename` command before you removed the node from the Solaris Volume Manager device group, run the `metaset` command from an active cluster node to remove the stale node entry from the Solaris Volume Manager device group. Then run the `clnode clear -F nodename` command to completely remove all traces of the node from the cluster.

Autodiscovery Should Find Only One Interconnect Path for Each Adapter (6299097)

Problem Summary: If there are redundant paths in the network hardware between interconnect adapters, the `scinstall` utility might fail to configure the interconnect path between them.

Workaround: If autodiscovery discovers multiple interconnect paths, manually specify the adapter pairs for each path.

Runtime

Logical Hostname Failover Could Create Duplicate Addresses, Lead To Outage (7201091)

Problem Summary: For a shared-IP zone-cluster (`ip-type=shared`), if the underlying non-global zone of a zone-cluster node is shut down by using the `uadmin 1 0` or `uadmin 2 0` command, the resulting failover of `LogicalHostname` resources might result in duplicate IP addresses being configured on a new primary node. The duplicate address is marked with the `DUPLICATE` flag until five minutes later, during which time the address is not usable by the application. See the `ifconfig(1M)` man page for more information about the `DUPLICATE` flag.

Workaround: Use either of the following methods:

- Cleanly shut down the zone-cluster node from the global zone.


```
# /usr/cluster/bin/clzonecluster -n nodename halt zone-cluster
```
- Before you perform any shutdown action from within the zone-cluster node, evacuate all resource groups from the zone-cluster node.


```
# /usr/cluster/bin/clresourcegroup evacuate -n zone-cluster-node +
```

sc_delegated_restarter Does Not Take Into Account Environment Variable Set in Manifest (7173159)

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

Workaround: There is no workaround.

Unable to Re-enable Transport Interface After Disabling With `ipadm disable-if -t interface` (7141828)

Problem Summary: Cluster transport paths go offline with accidental use of the `ipadm disable-if` command on the private transport interface.

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

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

```
# /usr/cluster/bin/clinterconnect show | grep Cable
```

2. Disable the cable for this interface on this node.

```
# /usr/cluster/bin/clinterconnect disable cable
```

3. Re-enable the cable to bring the path online.

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

Failure of Logical Hostname to Fail Over Caused by getnetmaskbyaddr() (7075347)

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

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

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

Upgrade

x86: scinstall -u update Sometimes Fails to Upgrade the Cluster Packages on an x86 Node (7201491)

Problem Summary: Running `scinstall -u update` on an x86 cluster node sometimes fails to upgrade the cluster packages. The following error messages are reported:

```
root@phys-schost-1:~# scinstall -u update
```

```
Calling "scinstall -u preupgrade"
```

```
Renamed "/.alt.s11u1_24a-2/etc/cluster/ccr" to "/.alt.s11u1_24a-2/etc/cluster/ccr.upgrade".
Log file - /.alt.s11u1_24a-2/var/cluster/logs/install/scinstall.upgrade.log.12037
```

```
** Upgrading software **
```

```
Startup: Linked image publisher check ... Done
Startup: Refreshing catalog 'aie' ... Done
Startup: Refreshing catalog 'solaris' ... Done
Startup: Refreshing catalog 'ha-cluster' ... Done
Startup: Refreshing catalog 'firstboot' ... Done
Startup: Checking that pkg(5) is up to date ... Done
```



```

Planning: Solver setup ... Done
Planning: Running solver ... Done
Planning: Finding local manifests ... Done
Planning: Fetching manifests: 0/26 0% complete
Planning: Fetching manifests: 26/26 100% complete
Planning: Package planning ... Done
Planning: Merging actions ... Done
Planning: Checking for conflicting actions ... Done
Planning: Consolidating action changes ... Done
Planning: Evaluating mediators ... Done
Planning: Planning completed in 16.30 seconds
Packages to update: 26

```

```

Planning: Linked images: 0/1 done; 1 working: zone:OtherNetZC
pkg: update failed (linked image exception(s)):

```

A 'update' operation failed for child 'zone:OtherNetZC' with an unexpected return value of 1 and generated the following output:
 pkg: 3/4 catalogs successfully updated:

```

Framework stall:
URL: 'http://bea100.us.oracle.com:24936/versions/0/'

```

Workaround: Before you run the `scinstall -u update` command, run `pkg refresh --full`.

Software Updates

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

- “Patch Management Tools” on page 47
- “My Oracle Support” on page 48

Note – Read the software update README before applying or removing any update.

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

For information about firmware updates, see your hardware documentation.

Patch Management Tools

Information for using the Oracle Solaris package management utility, `pkg`, is provided in Chapter 4, “Installing and Updating Software Packages,” in *Adding and Updating Oracle Solaris 11.1 Software Packages*.

My Oracle Support

The My Oracle Support Web site provides 24-hour access to the most up-to-date information regarding software, software updates, and firmware for Oracle products. Access the My Oracle Support site at <http://support.oracle.com> for the most current matrixes of supported software, firmware, and software update revisions.

Before you install Oracle Solaris Cluster 4.1 software and apply software updates to a cluster component (Oracle Solaris OS, Oracle Solaris Cluster software, volume manager software, data services software, or disk hardware), review each README file that accompanies the updates that you retrieved. All cluster nodes must have the same software update level for proper cluster operation.

For specific software update procedures, see [Chapter 11, “Updating Your Software,” in *Oracle Solaris Cluster System Administration Guide*](#).

Oracle Solaris Cluster 4.1 Documentation Set

The Oracle Solaris Cluster 4.1 user documentation is available in PDF and HTML format at the following web site:

http://docs.oracle.com/docs/cd/E29086_01/index.html

Documentation Issues

This section discusses errors or omissions for documentation in the Oracle Solaris Cluster and Geographic Edition 4.1 release.

- “Upgrade Guide” on page 48
- “HA for Oracle Guide” on page 49
- “HA for Oracle RAC Guide” on page 49
- “HA for Oracle Solaris Zones Guide” on page 50
- “Solaris Volume Manager” on page 50
- “Geographic Edition Data Replication Guide for Oracle Solaris Availability Suite” on page 50
- “Man Pages” on page 51

Upgrade Guide

In multiple chapters, the syntax for the `scinstall -u update` command is missing the option to specify license information, when needed. The full command syntax is the following:

```
# scinstall -u update -b [-b bename] [-L accept,licenses]
```


For more information about the `-L` option, see the `scinstall(1M)` man page.

HA for Oracle Guide

- In “Setting HA for Oracle Extension Properties” in *Oracle Solaris Cluster Data Service for Oracle Guide*, the list of required extension properties for the Oracle server resource is valid only if Oracle Grid Infrastructure is used. If you are not using Oracle Grid Infrastructure, the following extension properties are also required for the Oracle server resource:

- `Connect_string`
- `Alert_log_file`

This information is also missing from Step 9 of “How to Register and Configure HA for Oracle Without Oracle ASM (CLI)” in *Oracle Solaris Cluster Data Service for Oracle Guide*.

For information about the `Connect_string` and `Alert_log_file` extension properties, see the `SUNW.oracle_server(5)` man page.

- In “How to Prepare the Nodes” in *Oracle Solaris Cluster Data Service for Oracle Guide*, Step 7 is corrected and Step 8 is added as follows:

7. If you are using a zone cluster, configure the `limitpriv` property by using the `clzonecluster` command.

```
# clzonecluster configure zcname
clzonecluster:zcname>set limitpriv="default,proc_priocntl,proc_clock_highres"
clzonecluster:zcname>commit
```

8. On each zone-cluster node, prevent Oracle Clusterware time synchronization from running in active mode.

- Log in to the zone-cluster node as root.
- Create an empty `/etc/inet/ntp.conf` file.

```
# touch /etc/inet/ntp.conf
```

HA for Oracle RAC Guide

The procedure “How to Set the Necessary Privileges for Oracle RAC Software in a Zone Cluster” in Chapter 1 contains incorrect information. The correct procedure is as follows:

1. Become superuser on the global cluster node that hosts the zone cluster.
2. Configure the `limitpriv` property by using the `clzonecluster` command.

```
# clzonecluster configure zcname
clzonecluster:zcname>set limitpriv ="default,proc_priocntl,proc_clock_highres"
clzonecluster:zcname>commit
```

3. Beginning with Oracle RAC version 11g release 2, prevent Oracle Clusterware time synchronization from running in active mode.

- a. Log in to the zone-cluster node as root.
- b. Create an empty `/etc/inet/ntp.conf` file.

```
# touch /etc/inet/ntp.conf
```

HA for Oracle Solaris Zones Guide

The following instruction is missing from “[How to Install a Zone and Perform the Initial Internal Zone Configuration](#)” in *Oracle Solaris Cluster Data Service for Oracle Solaris Zones Guide*. Perform this step immediately after Step 6b:

- c. In the node where you updated the new UUID in the boot environment, if there are other non-global zones of brand type `solaris` configured, set the same UUID on the active boot environment for each non-global zone of brand type `solaris`.

```
phys-schost-2# zfs set org.opensolaris.libbe:parentbe=uuid poolname/zonepath/rpool/ROOT/bename
```

For example:

```
phys-schost-2# zoneadm list -cv
```

```
...
1 myzone1      running    /zones/myzone1  solaris  shared
...
```

```
phys-schost-2# zlogin myzone1 beadm list -H
```

```
solaris;4391e8aa-b8d2-6da9-a5aa-d8b3e6ed6d9b;NR;/;606941184;static;1342165571
```

```
phys-schost-2# zfs set org.opensolaris.libbe:parentbe=8fe53702-16c3-eb21-ed85-d19af92c6bbd \
rpool/zones/myzone1/rpool/ROOT/solaris
```

Solaris Volume Manager

Oracle Solaris Cluster 4.1 software supports Solaris Volume Manager software. The Oracle Solaris 11 documentation set does not include a manual for Solaris Volume Manager software. However, you can still use the *[Solaris Volume Manager Administration Guide](#)* from the Oracle Solaris 10 9/10 release, which is valid with the Oracle Solaris Cluster 4.1 release.

Geographic Edition Data Replication Guide for Oracle Solaris Availability Suite

The following instruction is missing from the procedure “How to Add an Application Resource Group to an Availability Suite Protection Group”:

If the application resource group to add is configured with a raw-disk device group, that device group must be specified in the resource group configuration by its data volume, rather than by its device group name. This ensures that the resource will remain monitored after the application resource group is added to a protection group.

For example, if the device group `rawd` has a corresponding data volume of `/dev/global/rdsk/d1s0`, you must set the `GlobalDevicePaths` property of the application resource group with the data volume, as follows:

```
# clresourcegroup set -p GlobalDevicePaths=/dev/global/rdsk/d1s0 rawdg
```

Man Pages

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

- [“clzonecluster\(1CL\)” on page 51](#)
- [“ORCL.sapcenter\(5\)” on page 51](#)
- [“ORCL.saprepenq\(5\)” on page 52](#)
- [“ORCL.saprepenq_preempt\(5\)” on page 52](#)
- [“ORCL.sapstartsrv\(5\)” on page 52](#)
- [“scdpmd.conf\(4\)” on page 52](#)
- [“scha_check_app_user\(1HA\)” on page 52](#)
- [“SUNW.HAStoragePlus\(5\)” on page 52](#)
- [“SUNW.ScalDeviceGroup\(5\)” on page 53](#)
- [“SUNW.ScalMountPoint\(5\)” on page 53](#)

clzonecluster(1CL)

- The (cluster) `ip`-type property incorrectly states that the only supported value is `shared`. Both `shared` and `exclusive` `ip`-types are supported.
- The `privnet` resource name incorrectly contains a hyphen (`priv-net`). The correct resource name is `privnet`.

ORCL.sapcenter(5)

- In the Description section, the seventh bullet point must read as follows:
 The resource group weak positive affinities must ensure that the SAP central service resource group fails over to the node where the SAP replicated enqueue resource group is online. If an `ORCL.saprepenq_preempt` resource is not configured, it must be implemented by strong negative affinities such that the replicated enqueue server resource group is off-loaded from the failover target node before the SAP central service resource group is started.
- In Example 1, make the following change:
 Change: `-p resource_dependencies=bono-1,db-rs,scs-strt-rs`
 To: `-p resource_dependencies=db-rs,scs-strt-rs`

ORCL.saprepenq(5)

In the Description section, the seventh bullet point must read as follows:

The resource group weak positive affinities must ensure that the SAP central service resource group fails over to the node where the SAP replicated enqueue resource group is online. If an `ORCL.saprepenq_preempt` resource is not configured, it must be implemented strong negative by affinities such that the replicated enqueue server resource group is off-loaded from the failover target node before the SAP central service resource group is started.

ORCL.saprepenq_preempt(5)

In the Description section, the eighth bullet point must read as follows:

The resource group weak positive affinities must ensure that the SAP central service resource group fails over to the node where the SAP replicated enqueue resource group is online. If an `ORCL.saprepenq_preempt` resource is not configured, it must be implemented by strong negative affinities such that the replicated enqueue server resource group is off-loaded from the failover target node before the SAP central service resource group is started. If the replicated enqueue preempt resource is configured, it is the task of this resource to off-load the replicated enqueue server resource group to a spare node after the enqueue tables are copied.

ORCL.sapstartsrv(5)

- In the Name section, the sentence describing the resource type must read as follows:
resource type implementation for processing `sapstartsrv` of Oracle Solaris Cluster HA for SAP NetWeaver
- In Example 1, make the following change:

```
Change: /usr/cluster/bin/clrs create -d -g pas-rg -t sapstartsrv  
To: /usr/cluster/bin/clrs create -d -g scs-rg -t sapstartsrv
```

scdpmd.conf(4)

The minimum value for the `Ping_interval` property is incorrect. The value should be 20, not 60.

scha_check_app_user(1HA)

The use of “effective user ID” in this man page is incorrect. The correct term in all places is “real user ID”. For information about the distinction between a real user ID and an effective user ID, see the [setuid\(2\)](#) man page.

SUNW.HAStoragePlus(5)

In the description for the `RebootOnFailure` property, the second paragraph is incorrect. The correct paragraph is the following:

If `RebootOnFailure` is set to `TRUE` and at least one device is found available for each entity specified in the `GlobalDevicePaths`, `FileSystemMountPoints`, or `Zpools` property, the local system is rebooted. The local system refers to the global-cluster node or the zone-cluster node where the resource is online.

SUNW.ScalDeviceGroup(5)

In the description for the `RebootOnFailure` property, the second paragraph is incorrect. The correct paragraph is the following:

If `RebootOnFailure` is set to `TRUE` and at least one device is found available for each entity specified in the `GlobalDevicePaths`, `FileSystemMountPoints`, or `Zpools` property, the local system is rebooted. The local system refers to the global-cluster node or the zone-cluster node where the resource is online.

SUNW.ScalMountPoint(5)

In the description for the `RebootOnFailure` property, the second paragraph is incorrect. The correct paragraph is the following:

If `RebootOnFailure` is set to `TRUE` and at least one device is found available for each entity specified in the `GlobalDevicePaths`, `FileSystemMountPoints`, or `Zpools` property, the local system is rebooted. The local system refers to the global-cluster node or the zone-cluster node where the resource is online.

ORCL.otd(5) Man Page

This appendix contains information about the `ORCL.otd` resource type, which is introduced in the Oracle Solaris Cluster data service for Oracle Traffic Director.

ORCL.otd(5) Man Page

`ORCL.otd`, `otd` - resource type implementation for Oracle Traffic Director

Description

The `ORCL.otd` resource type represents the Oracle Traffic Director application in an Oracle Solaris Cluster configuration.

The `ORCL.otd` resource type is a multi-master resource. It supports being configured in one of two modes. Although both modes allow the server instances to run on multiple nodes simultaneously, the difference is in the way the clients connect to the server instances. In the first mode, where logical hosts are used, the requests are received by the node that has the logical host resource online. In the second mode, where shared addresses are used, the load balancing feature of the Oracle Solaris Cluster software is used to distribute incoming client requests to the different instances. You can increase the utilization of the server instances by using multiple logical hosts.

The `ORCL.otd` resource type executes the following Oracle Traffic Director commands as the `userid` that owns the commands.

- `$ORACLE_HOME/lib/parsexml`
- `$INSTANCE_HOME/bin/startserv`
- `$INSTANCE_HOME/bin/stopserv`

The Oracle Traffic Director instance is started by the resource under the same `userid`. These commands are executable by the owner and are not writable by anyone except the owner.

You must set the following properties on an ORCL.otd resource by using `clresource(1CL)`.

Standard Properties

The standard resource properties, Scalable, Network_resources_used, Port_list, Load_balancing_policy, and Load_balancing_weights, are common to all scalable resource types. The properties controlling the behavior of the shared address load balancing feature, such as Scalable, Network_resources_used, Port_list, Load_balancing_policy, and Load_balancing_weights, are only applicable when the ORCL.otd resource type is configured with a shared address.

See [r_properties\(5\)](#) for a complete description of the following resource properties.

Load_balancing_policy

| | |
|------------------|-------------|
| Category | Optional |
| Data type | String |
| Default | Lb_weighted |
| Tunable | At creation |

Load_balancing_weights

| | |
|------------------|--------------|
| Category | Optional |
| Data type | String array |
| Default | NULL |
| Tunable | Any time |

Network_resources_used

| | |
|------------------|----------------------|
| Category | Conditional/Optional |
| Data type | String array |
| Default | Null |
| Tunable | At creation |

Resource_dependencies

| | |
|------------------|--------------|
| Category | Optional |
| Default | Null |
| Data type | String array |
| Tunable | Any time |

Port_list

Category Required
Data type String array
Default 80/tcp
Tunable When disabled

Retry_count

Category Conditional
Minimum 0
Maximum 10
Default 2
Tunable Any time

Retry_interval

Category Conditional
Minimum 0
Maximum 3600
Default 620
Tunable Any time

Thorough_probe_interval

Category Conditional
Minimum 0
Maximum 3600
Default 60
Tunable Any time

Extension Properties

Oracle_home

Category Required
Data type String
Default ---not set---

Tunable When disabled

This property is a per-node string of the Oracle Traffic Director installation home directory path name. You must specify this property for each node at resource creation time.

Instance_home

Category Required

Data type String

Default ---not set---

Tunable When disabled

This property is a per-node string of the Oracle Traffic Director instance directory path name. You must specify this property for each node at resource creation time.

Server_URL

Category Required

Data type String

Default Null

Tunable Any time

By default, the Oracle Traffic Director server instance is probed either through the `localhost:port` when a logical host is used, or the `shared_address:port` when a shared address is used.

This property allows you to specify an alternate URL to monitor the Oracle Traffic Director server instance. The fault monitor periodically runs the HTTP GET command for the URL specified and takes action if the HTTP request returns with response code 500 "Internal Server Error", or if the application server does not respond within the configured timeout period.

Monitor_retry_count

Category Required

Data type Integer

Minimum -1

Default 4

Tunable Any time

Controls the restarts of the fault monitor. This property indicates the number of times the fault monitor is restarted by the process monitor facility and corresponds to the `-n` option passed to the `pmfadm(1M)` command. The number of restarts is counted in a specified time

window. See the `Monitor_retry_interval` property for more information. Note that this property refers to the restarts of the fault monitor itself, not the web server. The restarts of the web server are controlled by the system-defined properties `Thorough_probe_interval`, `Retry_interval`, and `Retry_count`, as specified in their descriptions. See the [clresource\(1CL\)](#) man page for more information. You can modify the value for this property at any time.

`Monitor_retry_interval`

Category Required

Data type Integer

Minimum -1

Default 2

Tunable Any time

Indicates the time in minutes, over which the failures of the fault monitor are counted, and corresponds to the `-t` option passed to the `pmfadm(1M)` command. If the number of times the fault monitor fails exceeds the value of `Monitor_retry_count`, the fault monitor is not restarted by the Process Monitor Facility. You can modify the value for this property at any time.

`Probe_timeout`

Category Required

Data type Integer

Minimum 15

Default 90

Tunable Any time

This property is the time-out value, in seconds, that is used by the fault monitor to probe an Oracle Traffic Director instance. You can modify the value for this property at any time.

`Num_probes`

Category Required

Data type Integer

Minimum 1

Default 2

Tunable Any time

This property indicates the number of times that the fault monitor can timeout while probing before taking action on an Oracle Traffic Director server instance. You can modify the value for this property at any time.

Examples

The following assumptions are made for all the examples in this section.

The Oracle Traffic Director software is installed on a clustered file system in the `/global/otd/otd-home` directory as follows:

- The clustered file system is in a separate resource, `otd-gfs-rs`.
- The resource group is `otd-hasp-rg`.
- The mount point is `/global/otd`.

The Oracle Traffic Director server `Instance_home` directories are located on the same clustered file system as the installation, but each node has its own server `Instance_home` directory. For the examples in this section, the nodes and directories are as follows:

```
node1: /global/otd/otd-1/net-otd-a
node2: /global/otd/otd-2/net-otd-a
node3: /global/otd/otd-3/net-otd-a
node4: /global/otd/otd-4/net-otd-a
```

Use the following command to register the Oracle Traffic Director resource type:

```
# clresourcetype register ORCL.otd
```

EXAMPLE A-1 Creating an Oracle Traffic Director Resource for Use With a Logical Host

This example creates an Oracle Traffic Director resource, `otd-rs`, in a resource group, `otd-rg`. It is configured to run simultaneously on all the four nodes of a four-node cluster.

In an agent configuration where a logical host is being used, the Oracle Traffic Director server instances must be configured to listen on all addresses, `INADDR_ANY`, which allows the fault monitor to connect to the `localhost` address of each node using the default `Port_List`, `80/tcp`. The clients use the IP addresses as configured in a logical host resource, `lh-rs`, which is contained in the resource group, `lh-rg`. The hostname `otd-a-lh` is configured in the naming service used by the cluster and in any of the clients that will be accessing the server instances.

To create a logical host in this example:

```
# clresourcegroup create -p Nodelist="node1,node2,node3,node4"
-p Failback=True lh-rg \
# clreslogicalhostname create -g lh-rg -h otd-a-lh lh-rs \
# clresourcegroup online -eM lh-rg
```


EXAMPLE A-1 Creating an Oracle Traffic Director Resource for Use With a Logical Host *(Continued)*

To facilitate the automatic failover of the logical host to a node that has a running instance of Oracle Traffic Director:

- The logical host resource group must have a strong positive affinity with failover delegation to the Oracle Traffic Director resource group.
- The logical host resource must also have an offline-restart dependency on the Oracle Traffic Director resource with a local-node scope.

To create the Oracle Traffic Director resource group and resource, do the following:

```
# clresourcegroup create -S otd-rg \
# clresourcetype register ORCL.otd \
# clresource create -g otd-rg -t ORCL.otd \
-p ORACLE_HOME=/global/otd/otd-home \
-p INSTANCE_HOME{node1}=/global/otd/otd-1/net-otd-a \
-p INSTANCE_HOME{node2}=/global/otd/otd-2/net-otd-a \
-p INSTANCE_HOME{node3}=/global/otd/otd-3/net-otd-a \
-p INSTANCE_HOME{node4}=/global/otd/otd-4/net-otd-a \
-p Resource_dependencies_offline_restart=otd-gfs-rs \
-p Port_List=80/tcp otd-rs \
# clresourcegroup set -p RG_affinities++++otd-rg lh-rg \
# clresource set -p Resource_dependencies_offline_restart+=otd-rs{local_node} lh-rs
```

EXAMPLE A-2 Creating an Oracle Traffic Director Resource for Use With a Shared Address

This example creates an Oracle Traffic Director otd-rs resource named otd-rg in a resource group named web-rg, which is configured to run simultaneously on all four nodes of a four-node cluster.

The Oracle Traffic Director server instances are configured to listen on port 80 and uses the IP addresses as configured in a SharedAddress resource named sa-rs, which is contained in the resource group sa-rg. The hostname otd-a-sa, is configured in the naming service used by the cluster and any of the clients that will be accessing the server instances.

To create the shared address resource group and resource for this example, do the following:

```
# clresourcegroup create sa-rg \
# clressharedaddress create -g sa-rg -h otd-a-sa sa-rs \
# clresourcegroup online -eM sa-rg
```

To create the Oracle Traffic Director resource group and resource, do the following:

```
# clresourcegroup create -S otd-rg \
# clresourcetype register ORCL.otd \
# clresource create -g otd-rg -t ORCL.otd \
-p ORACLE_HOME=/global/otd/otd-home \
-p INSTANCE_HOME{node1}=/global/otd/otd-1/net-otd-a \
-p INSTANCE_HOME{node2}=/global/otd/otd-2/net-otd-a \
```


EXAMPLE A-2 Creating an Oracle Traffic Director Resource for Use With a Shared Address
(Continued)

```
-p INSTANCE_HOME{node3}=global/otd/otd-3/net-otd-a \  
-p INSTANCE_HOME{node4}=global/otd/otd-4/net-otd-a \  
-p Resource_dependencies_offline_restart=otd-gfs-rs \  
-p Resource_dependencies=sa-rs \  
-p Port_List=80/tcp \  
-p Scalable=True \  
otd-rs
```

EXAMPLE A-3 Setting up an Alternate Monitoring Server_URL

To configure the agent fault monitor to probe a specific URL, specify the URL to be probed in the Server_URL extension property.

The following example shows how to set the Server_URL extension property. In this case, the Oracle Traffic Director server instances are configured to have the statistics monitoring enabled and set to the `http://localhost:80/stats-xml` URL.

```
# clresource set -p server_url="http://localhost:80/stats-xml" \  
otd-rs
```

Attributes

See [attributes\(5\)](#) for description of the following attribute:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------------------------|
| Availability | ha-cluster/data-service/ORCLscotd |

See Also

[clresource\(1CL\)](#), [clressharedaddress\(1CL\)](#), [clreslogicalhostname\(1CL\)](#), [clresourcetype\(1CL\)](#), [clresourcegroup\(1CL\)](#), [pmfadm\(1M\)](#), [attributes\(5\)](#), [r_properties\(5\)](#), [scalable_service\(5\)](#)

Oracle Solaris Cluster Data Services Planning and Administration Guide