Oracle® Solaris Cluster 4.0 Release Notes

Part No: E23438
March 2013, E23438-06

ORACLE

130327@25097

Contents

Preface

This document lists product features, requirements, and restrictions for the Oracle Solaris Cluster 4.0 and Oracle Solaris Cluster Geographic Edition (Geographic Edition) 4.0 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 `.login` file. Use `ls -a` to list all files. `machine_name% you have mail.`
`AaBbCc123`	What you type, contrasted with onscreen computer output	`machine_name%` **`su`** `Password:`
aabbcc123	Placeholder: replace with a real name or value	The command to remove a file is `rm` *filename*.
AaBbCc123	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the *User's Guide*. A *cache* is a copy that is stored locally. Do *not* save the file. **Note:** Some emphasized items appear bold online.

Shell Prompts in Command Examples

The following table shows the default UNIX system prompt and superuser prompt for shells that are included in the Oracle Solaris OS. Note that the default system prompt that is displayed in command examples varies, depending on the Oracle Solaris release.

TABLE P–2 Shell Prompts

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

Related Documentation

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

Topic	Documentation
Hardware installation and administration	*Oracle Solaris Cluster 4.0 Hardware Administration Manual* Individual hardware administration guides
Concepts	*Oracle Solaris Cluster Concepts Guide*
Software installation	*Oracle Solaris Cluster Software Installation Guide*
Data service installation and administration	*Oracle Solaris Cluster Data Services Planning and Administration Guide* and individual data service guides
Data service development	*Oracle Solaris Cluster Data Services Developer's Guide*
System administration	*Oracle Solaris Cluster System Administration Guide* *Oracle Solaris Cluster Quick Reference*
Software upgrade	*Oracle Solaris Cluster Upgrade Guide*
Error messages	*Oracle Solaris Cluster Error Messages Guide*
Command and function references	*Oracle Solaris Cluster Reference Manual* *Oracle Solaris Cluster Data Services Reference Manual* *Oracle Solaris Cluster Geographic Edition Reference Manual* *Oracle Solaris Cluster Quorum Server Reference Manual*

For a complete list of Oracle Solaris Cluster documentation, see the release notes for your version of Oracle Solaris Cluster software.

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 4.0)

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

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

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

Oracle Solaris Cluster 4.0 Release Notes

This document provides the following information for the Oracle Solaris Cluster 4.0 software and the Oracle Solaris Cluster Geographic Edition 4.0 software.

What's New in the Software

This section provides information related to new features, functionality, and products in the Oracle Solaris Cluster 4.0 and Oracle Solaris Cluster Geographic Edition (Geographic Edition) 4.0 software.

- Oracle Solaris Cluster is an integrated hardware and software solution that helps you create highly available and scalable services. Geographic Edition is a layered extension of the Oracle Solaris Cluster software that protects applications from unexpected disruptions by using multiple clusters that are separated by long distances. Geographic Edition also uses a redundant infrastructure that replicates data between these clusters.

The Oracle Solaris Cluster 4.0 software provides the following new features:

Automated Installer Support

The Oracle Solaris Cluster 4.0 software supports the use of the Oracle Solaris 11 Automated Installer (AI) to install cluster software. AI is similar to the Oracle Solaris 10 JumpStart feature, which was used in the Oracle Solaris Cluster 3.3 release as a method of installing packages. You can use AI for a hands-free installation of the packages for the Oracle Solaris Cluster 4.0 software.

For information about AI, see the Part III, "Installing Using an Install Server," in *Installing Oracle Solaris 11 Systems*. For more information on installing and uninstalling the Oracle Solaris Cluster 4.0 software, see the *Oracle Solaris Cluster Software Installation Guide*.

New Cluster Package Names

Oracle Solaris Cluster packages have been renamed to new Oracle Solaris Image Packaging System (IPS) package names. The `pkgadd` and `pkgrm` commands have been replaced with IPS commands, such as `pkg install`.

This section provides the following lists of individual IPS packages for the Oracle Solaris Cluster 4.0 release that had SVR4 equivalents in previous Oracle Solaris Cluster releases:

- Table 1
- Table 2
- Table 3

These lists are not comprehensive; they do not include group packages or individual IPS packages that had no SVR4 equivalent. You can view the complete list of Oracle Solaris Cluster 4.0 packages and package groups from the package repository.

TABLE 1 New IPS Package Names for SPARC and x86

Previous Package Name	New IPS Package Name
SUNWscapc	ha-cluster/data-service/apache
SUNWscdhc	ha-cluster/data-service/dhcp
SUNWscdns	ha-cluster/data-service/dns
SUNWscs1as	ha-cluster/data-service/glassfish-app-server

TABLE 1 New IPS Package Names for SPARC and x86 *(Continued)*

Previous Package Name	New IPS Package Name
SUNWscs1mq	ha-cluster/data-service/glassfish-message-queue
SUNWsczone	ha-cluster/data-service/ha-zones
SUNWscnfs	ha-cluster/data-service/nfs
SUNWscor	ha-cluster/data-service/oracle-database
SUNWsctomcat	ha-cluster/data-service/tomcat
SUNWscwls	ha-cluster/data-service/weblogic
SUNWscdsbuilder	ha-cluster/developer/agent-builder
SUNWscdev	ha-cluster/developer/api
SUNWscderby	ha-cluster/ha-service/derby
SUNWscgds	ha-cluster/ha-service/gds
SUNWscrtlh	ha-cluster/ha-service/logical-hostname
SUNWscsmf	ha-cluster/ha-service/smf-proxy
SUNWsctelemetry	ha-cluster/ha-service/telemetry
SUNWsccacao	ha-cluster/library/cacao
SUNWscucm	ha-cluster/library/ucmm
SUNWesc, SUNWfsc, SUNWjsc, SUNWcsc	ha-cluster/locale
SUNWscnmr, SUNWscnmu	ha-cluster/release/name
SUNWscmasar, SUNWscmasazu, SUNWscmautil, SUNWscmautilr	ha-cluster/service/management
SUNWscmasasen	ha-cluster/service/management/slm
SUNWscqsr, SUNWscqsu	ha-cluster/service/quorum-server
SUNWscqsman	ha-cluster/service/quorum-server/manual
SUNWjscqsu, SUNWcscqsu	ha-cluster/service/quorum-server/locale
SUNWjscqsman	ha-cluster/service/quorum-server/manual/locale
SUNWmdmr, SUNWmdmu	ha-cluster/storage/svm-mediator
SUNWscsckr, SUNWscscku	ha-cluster/system/cfgchk

TABLE 1 New IPS Package Names for SPARC and x86 *(Continued)*

Previous Package Name	New IPS Package Name
SUNWsc, SUNWscu, SUNWscr, SUNWsczr, SUNWsczu, SUNWsccomu, SUNWsccomzu	ha-cluster/system/core
SUNWscmasa, SUNWscmasau	ha-cluster/system/dsconfig-wizard
SUNWscman	ha-cluster/system/manual
SUNWscdsman	ha-cluster/system/manual/data-services
SUNWjscman	ha-cluster/system/manual/locale

TABLE 2 New IPS Package Names for SPARC only

Previous Package Name	New IPS Package Name
SUNWscxvm	ha-cluster/data-service/ha-ldom

TABLE 3 New IPS Package Names for Geographic Edition

Previous Geographic Edition Package Name	New IPS Package Name
SUNWscgctl, SUNWscgctlr, SUNWscghb, SUNWscghbr	ha-cluster/geo/framework
SUNWscgrepavs, SUNWscgrepavsu	ha-cluster/geo/replication/availability-suite
SUNWscgrepodg, SUNWscgrepodgu	ha-cluster/geo/replication/data-guard
SUNWscgrepsbpu	ha-cluster/geo/replication/sbp
SUNWscgman	ha-cluster/geo/manual

Default Root File System of Oracle Solaris ZFS

The Oracle Solaris ZFS file system is now the default root file system for Oracle Solaris 11 and works with the Oracle Solaris Cluster 4.0 software.

Selected Support for Non-Global Zones

The Oracle Solaris Cluster 4.0 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.

HA for Oracle with Oracle Data Guard Replication

Support for Oracle Data Guard replication is expanded to include highly available (HA) Oracle database deployments. The minimum Oracle version for this support is 10.2.0.4.3.

This new feature also enables physical standby databases to be integrated into a Oracle Solaris Cluster deployment. Possible configurations include using HA for Oracle at both the primary and standby sites, or using a mix of HA for Oracle and Oracle RAC.

As of the time of publication, logical-standby configurations have not been tested and are not supported.

For more information, see Oracle Solaris Cluster Geographic Edition Data Replication Guide for Oracle Data Guard.

What's Not Included in the Oracle Solaris Cluster 4.0 Software

The following features were included in the Oracle Solaris Cluster 3.3 release but are not included in the Oracle Solaris Cluster 4.0 release:

- Support for Veritas File System (VxFS) and Veritas Volume Manager (VxVM)
- Support for the VxVM cluster feature for Oracle RAC in addition to VxVM with Oracle Solaris Cluster
- Support for 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)
- Support for Oracle Solaris Trusted Extensions
- The `scsnapshot` tool
- The `cconsole` utility (the Oracle Solaris `pconsole` utility can be used instead)

The following HA data services were not initially available with the 4.0 release but might become available at a later time:

- Afga IMPAX
- ASE

- Informix
- Kerberos
- MySQL
- Oracle Business Intelligence Enterprise Edition
- Oracle eBusiness Suite
- Oracle iPlanet Web Server
- PeopleSoft Enterprise
- PostgreSQL
- Samba
- SAP
- SAP liveCache
- SAP Web Application Server
- Siebel, SWIFTAlliance Access and Gateway
- Sybase
- TimesTen
- WebSphere Message Broker
- WebSphere Message Queue

The Grid Engine and Sun Java System Application Server EE (formerly called HADB) data services have been removed from the Oracle Solaris Cluster software.

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 4.0 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 Oracle support services to see if a code fix becomes 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, the error message `ORACLE_SID (+ASM2) does not match the Oracle ASM configuration ORACLE_SID () within CRS` or `ERROR: Oracle ASM is either not installed or the installation is invalid!` is reported by the `clsetup` utility. This situation occurs because, after Oracle Grid Infrastructure 11.2.0.3 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
```

IP Addresses on a Failed IP Interface Can No Longer Be Used Locally (7099852)

Problem Summary: This problem affects data services that use the `connect()` call to probe the health of the application through its logical hostname IP address. In a cluster-wide network outage scenario, there is a change in the behavior of the `connect()` call on the Oracle Solaris 11 software from the Oracle Solaris 10 release. The `connect()` call fails if the IPMP interface, on which the logical hostname IP is plumbed, goes down. This makes the agent probe fail if the network outage is longer than the `probe_timeout` and eventually brings the resource and the associated resource group to the offline state.

Workaround: Configure the application to listen on *localhost*:*port* to ensure that the monitoring program does not fail the resource in a public-network outage scenario.

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.

DID Disk Add to Solaris Zone Is Not Accepting Wild Card for *dsk (7081090)

Problem Summary: When using the `zonecfg` utility, if you add a DID disk to a non-global zone by using a wild card (*) and without specifying the paths, the addition fails.

Workaround: Specify the raw device paths and block device paths explicitly. The following example adds the `d5` DID device:

```
root@phys-cluster-1:~# zonecfg -z foo
zonecfg:foo> add device
zonecfg:foo:device> set match=/dev/did/dsk/d5s*
zonecfg:foo:device> end
zonecfg:foo> add device
zonecfg:foo:device> set match=/dev/did/rdsk/d5s*
zonecfg:foo:device> end
zonecfg:foo> exit
```

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 4.0 and Oracle Solaris Cluster Geographic Edition 4.0 software.

Data Replication

The following data replication products are supported in this release:

- **Availability Suite feature of Oracle Solaris** – When used with Geographic Edition software, support requires a minimum of Oracle Solaris 11 software repository update (SRU) 1.
- **Oracle Data Guard** – Support is limited to configurations that use Oracle Database 11.2.0.3 software.
- **Oracle Solaris Cluster Geographic Edition script-based plug-ins**.

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
Oracle Data Guard[1]	Oracle Automatic Storage Management
	Oracle Solaris ZFS Volume Manager
	Solaris Volume Manager for Sun Cluster
Availability Suite feature of Oracle Solaris	Solaris 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

Data services for the following applications are supported with Oracle Solaris Cluster 4.0:

- Apache
- Apache Tomcat

- DHCP
- DNS
- NFS
- Oracle GlassFish Server Message Queue
- Oracle Database single-instance on version 11.2.0.3 only
- Oracle Real Application Clusters on version 11.2.0.3 only
- Oracle Solaris Zones
- Oracle VM Server for SPARC
- Oracle WebTier
- WebLogic Server

The Grid Engine and Sun Java System Application Server EE (formerly called HADB) Data Services are no longer supported.

File Systems

The following table describes the file systems that work with the Oracle Solaris Cluster 4.0 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)	UFS and raw disk are the only possible cluster file systems

Oracle Solaris Cluster Geographic Edition Software Requirements

The Oracle Solaris Cluster Geographic Edition (Geographic Edition) 4.0 software runs on only Oracle Solaris Cluster 4.0 software.

A Geographic Edition configuration can consist of one cluster running Geographic Edition 4.0 software and the other cluster running Geographic Edition 3.3 5/11 software.

Memory Requirements

The Oracle Solaris Cluster 4.0 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 4.0, Oracle Solaris Cluster Geographic Edition 4.0, and Quorum Server software require the Oracle Solaris 11 OS.

Oracle VM Server for SPARC

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

Volume Management Software

The following table describes the volume managers that work with Oracle Solaris 11 on the SPARC and x86 platforms. Veritas Volume Manager (VxVM) is not supported for the Oracle Solaris Cluster and Oracle Solaris Cluster Geographic Edition 4.0 software.

Volume Manager	Cluster Feature
Solaris Volume Manager	Solaris Volume Manager for Sun Cluster
Oracle Solaris ZFS Volume Manager	

Product Localization

The following table describes localization for certain components of Oracle Solaris Cluster 4.0 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	`$ LC_MESSAGES=C;export LC_MESSAGES`
ksh	`$ export LC_MESSAGES=C`
bash	`$ export LC_MESSAGES=C`
csh	`% setenv LC_MESSAGES C`
tcsh	`% setenv LC_MESSAGES C`

Known Issues and Bugs

The following known issues and bugs affect the operation of the Oracle Solaris Cluster and Oracle Solaris Cluster Geographic Edition 4.0 software, as of the time of release. Bugs and issues are grouped into the following categories:

Check with Oracle support services to see if a code fix becomes available.

Administration

x86: `clzonecluster export` Command Fails (7066586)

Problem Summary: The following command might fail on x86 machines.

```
# clzonecluster export zonename
usage:
export [-f output-file]
```

Workaround: Use the following command instead:

```
# zonecfg -z zone-cluster-name export
```

Using `chmod` to `setuid` Returns Error in Non–Global Zone on PxFS Secondary Server (7020380)

Problem Summary: The `chmod` command run from a non-global zone might fail on a cluster file system. The `chmod` operation succeeds from a non-global zone on a node where the PxFS primary is located but fails from a non-global zone on a node where the PxFS secondary is located. For example:

```
# chmod 4755 /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.

`Cluster.CCR: libpnm system error: Failed to resolve pnm proxy pnm_server.2.`*zonename* (6942090)

Problem Summary: If using `solaris10` branded non-global zones with exclusive IP on an Oracle Solaris Cluster host, the `clnode status` command with the `-m` or `-v` option reports an error in the `/var/adm/messages` file similar to the following:

```
Cluster.CCR: [ID 544775 daemon.error] libpnm system error: Failed to resolve pnm proxy zonename
```

This error does not affect the running of the non-global zone or the cluster. The `solaris10` branded zone does not have to be under cluster control for the errors to be seen.

The issue is only seen on `solaris10` branded zones with exclusive IP. The issue is not seen when the following conditions exist:

- The zone is `solaris` branded and uses exclusive IP.
- The zone is `solaris10` branded and uses shared IP.
- The zone is `solaris` branded and uses shared IP.

Workaround: There is no workaround. The error messages do not affect the running of the non-global zone or the global cluster.

Missing `/dev/rmt` Causes Incorrect Reservation Usage When Policy Is pathcount (6920996)

Problem Summary: When a new storage device is added to a cluster and is configured with three or more DID paths, the node on which the `cldevice populate` command is run might fail to register its PGR key on the device.

Workaround: Run the `cldevice populate` command on all cluster nodes, or run the `cldevice populate` command twice from the same node.

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

Public 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.

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.

Data Services

Share Mount Point Matching Is Incorrect for Combination of UFS and ZFS Starting With a Common Pattern (7093237)

Problem Summary: If an NFS resource is created for a ZFS mount point and this mount-point prefix matches a UFS file system entry in the `vfstab` file, the HA for NFS data service will fail validation if the UFS file system is not mounted on the node.

Workaround: Mount the UFS file system on the node where the HAStoragePlus resource pertaining to the ZFS file system is online. You only need to do this if the resource is being created or updated. At any other time, there is no constraint that the UFS file system must be mounted before the resource group can be taken offline or brought online and the resource group can be switched to any node at will.

'Unable to Determine Oracle CRS Version' Error After Applying Patch 145333-09 (7090390)

Problem Summary: The Oracle Solaris Cluster code is unable to determine the Oracle CRS version when the `su` user is using the `csh` shell.

Workaround: A user that owns `${CRS_HOME}/bin/srvctl` must not use the `csh` shell.

SPARC: HA for Oracle VM Server for SPARC Default STOP_TIMEOUT is Too Low - Need Better Monitoring Of Domain Migration Progress (7069269)

Problem Summary: The `STOP_TIMEOUT` value in the HA for Oracle VM Server for SPARC data service is too low to complete the migration of guest domains.

Workaround: Increase the default value for `STOP_TIMEOUT` to at least 900 or to the expected migration time interval multiplied by 4.

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.

Installation

`cluster check` Fails for `cacaoadm` With `Insufficient Data` Before Node Is Configured in Cluster (7104375)

Problem Summary: The `cluster check` command uses common agent container (CAC) services for communication between nodes and requires CAC to be running. If any administrator runs the check S6979686 while the node is not a cluster member and CAC services are not running, the following message is displayed:

```
Insufficient Data: 1; /usr/sbin/cacaoadm status: Unable to check
SMF status
```

Workaround: This error is safe to ignore. Select the option `ignore the error and continue` to continue while installing the Oracle Solaris Cluster software.

Some Cluster Services Might Be Missing After Configuring Cluster on a Boot Environment That Previously Had the Cluster Software Installed (7103721)

Problem Summary: If you uninstall Oracle Solaris Cluster and then reinstall and configure it in the same boot environment, the cluster will boot successfully, but some of the cluster services might be missing. Run the `svcs -x` command and check for any services beginning with `svc:/system/cluster`.

```
# svcs -x
svc:/system/cluster/rgm-starter:default (Resource Group Manager Daemon)
 State: offline since Fri Oct 28 18:30:36 2011
Reason: Dependency svc:/system/cluster/rpc-fed:default is absent.
   See: http://sun.com/msg/SMF-8000-E2
Impact: 5 dependent services are not running.  (Use -v for list.)
```

Workaround: Use the following commands to add the absent service. The following example shows the addition of the `svc:/system/cluster/rpc-fed:default` service:

```
# service=svc:/system/cluster/rpc-fed:default
# svccfg -s ${service%:*} add ${service##*:}
# svccfg -s ${service} addpg general framework
# svccfg -s ${service} delcust -M
# svcadm enable ${service}
```

Then rerun the `svcs -x` command to check for any other missing cluster services.

`scinstall` Tries to Create an IPMP Group on a Standby Interface (7095759)

Problem Summary: If the cluster nodes have IPMP groups created with an active-standby configuration before Oracle Solaris Cluster configuration is performed, the `scinstall` command will fail with the following error messages during Oracle Solaris Cluster configuration:

```
Configuring IP multipathing groups ...failed
scinstall: Failed to retrieve the broadcast value for this adapter
```

If the standby adapter does not have any broadcast value, the `scinstall` command prints the above error message and does not proceed further in group creation. The `scinstall` command will, however, continue further without any issues.

Workaround: No workaround is required and the message is safe to ignore.

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

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 `netmasks`name service. This call to `getnetmaskbyaddr()` hangs for a while due to CR 7051511, which might hang long enough for the Resource Group Manager (RGM) to put the resource in the FAILED state. This occurs even though the correct netmask entries are in the `/etc/netmasks` local files. This issue affects only multi-homed clusters, such as cluster nodes that reside on multiple subnets.

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

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

ssm_start Fails Due to Unrelated IPMP Down (6938555)

Problem Summary: A scalable resource that depends on a `SUNW.SharedAddress` resource fails to come online, due to failure of an IPMP group that is on a subnet that is not used by the shared-address resource. Messages similar to the following are seen in the syslog of the cluster nodes:

```
Mar 22 12:37:51 schost1 SC SUNW.gds:5,Traffic_voip373,Scal_service_voip373,SSM_START:
ID 639855 daemon.error IPMP group sc_ipmp1 has status DOWN. Assuming this
node cannot respond to client requests.
```

Workaround: Repair the failed IPMP group and restart the failed scalable resource.

Software Updates

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

- "Patch Management Tools" on page 26
- "My Oracle Support" on page 26

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 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 4.0 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 4.0 Documentation Set

The Oracle Solaris Cluster 4.0 user documentation is available in PDF and HTML format at the following web site:

`http://docs.oracle.com/docs/cd/E23623_01/index.html`

Documentation Issues

This section discusses errors or omissions for documentation in the Oracle Solaris Cluster and Geographic Edition 4.0 release.

HA for Zones Procedure Moved to the Data Service Manual

The initial version of this Release Notes contained the procedure "How to Configure the HA for Zones Zone Boot Component for `solaris` or `solaris10` Brand Zones". That procedure was removed in an update of this Release Notes and can now be found at "How to Create and Enable Resources for the Zone Boot Component" in *Oracle Solaris Cluster Data Service for Oracle Solaris Zones Guide*.

Solaris Volume Manager

Oracle Solaris Cluster 4.0 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 4.0 release.

Man Pages

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

Section 3HA Man Pages

If you are developing an agent for services that will run in a zone cluster; and if your agent might need to execute some of its methods in the global zone; you can refer to the Oracle Solaris Cluster 3.3 5/11 version of the section 3HA man pages for information that was inadvertently omitted or altered in the 4.0 version of the section 3HA man pages.

clconfiguration(5)

The `globaldevfs` property is no longer valid and should be ignored.

clnasdevice(1CL)

At time of initial release, no NAS devices of type `sun` or `netapp_nas` are available. Information about the `sun` or `netapp_nas` NAS device type should be ignored.

clnode(1CL)

The description for the `remove` subcommand includes the following statement:

> This subcommand also removes the cluster software from the node.

This statement is incorrect and should be ignored. You must use the `pkg remove` command to remove the cluster software packages from a node.

clquorum(1CL)

At time of initial release, no Sun Microsystems, Inc. or Network Appliance (NetApp) NAS devices are available. Information about these NAS devices should be ignored.

cluster(1CL)

Example output has entries that mention the `Pkglist` property. This property is not used in the 4.0 release and the example content should be ignored.

clzonecluster(1CL)

-c *config_profile*.xml Option

The `-c` *config_profile*`.xml` option is added to the `install` subcommand. The following is the command syntax for this option.

```
clzonecluster install -c config_profile.xml zone-cluster-name
```

`-c` *config_profile*`.xml`
`--configprofile` *config_profile*`.xml`
Specifies a configuration profile template. After installation from the repository, the template applies the system configuration information to all nodes of the zone cluster. If *config_profile*`.xml` is not specified, you must manually configure each zone-cluster node by running from the global zone on each node the `zlogin -C` *zone-cluster-name* command. All profiles must have a `.xml` extension.

The `-c` option replaces the hostname of the zone-cluster node in the configuration profile template. The profile is applied to the zone-cluster node after booting the zone-cluster node.

Correction to Default Set of Packages That Are Installed by the Automated Installer

In the description of the `install` subcommand, the man page incorrectly states that, if you do not specify the `-M` option, the Automated Installer installs the `ha-cluster-full` group package by default. Instead, when `-M` is not specified, all of the `ha-cluster/*` packages that are installed in the global zone of the issuing node are installed in all nodes of the zone cluster.

Missing Description of the export Subcommand

The following syntax and description for the `export` subcommand is missing from the man page:

`/usr/cluster/bin/clzonecluster export [-f` *commandfile*`] zoneclustername`

`export` Exports the zone cluster configuration into a command file.

The exported *commandfile* can be used as the input for the `configure` subcommand. You can use the `export` subcommand only from a global-cluster node.

The RBAC authorization for the `export` subcommand is `solaris.cluster.admin`.

r_properties(5)

The following information applies to the `r_properties`(5) man page.

- Multiple instances of `Global_zone_override` were changed to `_override`.
- The `Resource_project_name` property description was omitted. Refer to the Oracle Solaris Cluster 3.3 5/11 version of the `r_properties`(5) man page for information about the `Resource_project_name` property.

- If you are developing an agent for services that will run in a zone cluster; and if your agent might need to execute some of its methods in the global zone; then you should refer to the Oracle Solaris Cluster 3.3 5/11 version of the `r_properties`(5) man page for information that was inadvertently omitted or altered in the 4.0 version of the `r_properties`(5) man page.

scinstall(1M)

The `-L` option is omitted from the `scinstall`(1M) man page. This option is used with the `scinstall -u update` command. The following is the syntax for specifying the `-L` option:

```
scinstall -u update [-b beename] [-L {accept | licenses | accept,licenses | licenses,accept}]
```

The argument `accept` corresponds to the `--accept` option of the `pkg` command and the argument `licenses` corresponds to the `--licenses` option.

Specifying `-L accept` indicates that you agree to and accept the licenses of the packages that are updated. If you do not provide this option, and any package licenses require acceptance, the update operation fails.

Specifying `-L licenses` displays all of the licenses for the packages that are updated.

When both `-L accept` and `-L licenses` are used, the licenses of the packages that are updated are displayed as well as accepted. The order you specify the `accept` and `licenses` arguments does not affect the behavior of the command.

rt_properties(5)

If you are developing an agent for services that will run in a zone cluster; and if your agent might need to execute some of its methods in the global zone; then you should refer to the Oracle Solaris Cluster 3.3 5/11 version of the `rt_properties`(5) man page for information that was inadvertently omitted or altered in the 4.0 version of the `rt_properties`(5) man page.

SUNW.gds(5)

The following extension properties are missing from the `SUNW.gds`(5) man page.

`Monitor_retry_count`

The number of times that the process monitor facility (PMF) restarts the fault monitor during the time window that the `Monitor_retry_interval` property specifies. This property refers to restarts of the fault monitor itself rather than to the resource. The system-defined properties `Retry_interval` and `Retry_count` control restarting of the resource.

Category Optional

Data type Integer

Default 4

Range 0 - 2147483647

-1 indicates an infinite number of retry attempts.

Tunable At any time

`Monitor_retry_interval` The time (in minutes) over which failures of the fault monitor are counted. If the number of times that the fault monitor fails exceeds the value that is specified in the extension property `Monitor_retry_count` within this period, the PMF does not restart the fault monitor.

Category Optional

Data type Integer

Default 2

Range 0 – 2147483647

-1 indicates an infinite retry interval.

Tunable At any time

SUNW.oracle_server(5)

The following value for the `Standby_mode` extension property is missing from the man page:

`SNAPSHOT` Beginning with Oracle 11*g*, specifies a snapshot standby database.

◆ ◆ ◆

ORCL.ohs(5) and ORCL.opmn(5) Man Pages

This appendix contains information about the `ORCL.ohs` and `ORCL.opmn` resource types, which are introduced in the Oracle Solaris Cluster data service for Oracle Web Tier.

ORCL.ohs(5) Man Page

`ORCL.ohs`, `ohs` - resource type implementation for failover Oracle HTTP Server (OHS)

Description

The `ORCL.ohs` resource type represents one of the HA for Oracle Web Tier components available to an Oracle Solaris Cluster configuration. The HA for Oracle HTTP Server resource can be configured only as a failover service.

The `ORCL.ohs` resource type is derived from, and extends, the `SUNW.gds` resource type. As a consequence the `Probe_command`, `Start_command`, `Stop_command`, and `Validate_command` properties are not tunable. However, all the `SUNW.gds` standard and extension resource type properties can be tuned. For a list of the `SUNW.gds` standard and extension properties, see the `SUNW.gds`(5) man page.

The `ORCL.ohs` resource type defines the following extension properties, in addition to the `SUNW.gds` extension properties that the resource type uses.

Extension Properties

`Component_instance`
The name of the Oracle HTTP Server instance that is listed in the `opmnctl` output.

Category	Required (if not specified, the default value is used)
Per Node	False
Data Type	String
Default	`ohs1`
Tunable	When disabled

`Debug_level`
The `Debug_level` property determines how much debugging information is produced during resource creation, update and during probe cycles.

Category	Optional
Per Node	True
Data Type	Integer
Minimum	`0`
Maximum	`2`
Default	`0`
Tunable	Any time

Attributes

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	`ha-cluster/data-service/ohs`

See Also

`pmfadm(1M)` `scha_resource_get(3HA)`, `clresourcetype(1CL)`, `attributes(5)`, `ORCL.opmn(5)`, `r_properties(5)`, `SUNW.gds(5)`

Oracle Solaris Cluster Data Services Planning and Administration Guide, *Oracle Solaris Cluster Data Service for Oracle Web Tier Guide*

ORCL.opmn(5) Man Page

`ORCL.opmn`, `opmn` - resource type implementation for failover Oracle Process Management and Notification Server (OPMN)

Description

The `ORCL.opmn` resource type represents one of the HA for Oracle Web Tier components available to an Oracle Solaris Cluster configuration. The HA for Oracle HTTP Server resource can be configured only as a failover service.

The `ORCL.opmn` resource type is derived from, and extends, the `SUNW.gds` resource type. As a consequence the `Probe_command`, `Start_command`, `Stop_command`, and `Validate_command` properties are not tunable. However, all the `SUNW.gds` standard and extension resource type properties can be tuned. For a list of the `SUNW.gds` standard and extension properties, see the `SUNW.gds(5)` man page.

The `ORCL.opmn` resource type defines the following extension properties, in addition to the `SUNW.gds` extension properties that the resource type uses.

Extension Properties

`Debug_level`
: The `Debug_level` property determines how much debugging information is produced during resource creation, updates, and probe cycles.

Category	Optional
Per Node	True
Data Type	Integer
Minimum	`0`
Maximum	`2`
Default	`0`
Tunable	Any time

`Instance_name`
: The instance name. A directory of this name must exist within the `ORACLE_HOME/instances` subdirectory.

Category	Required (if not specified, the default value is used)
Per Node	False
Data Type	String

Default `instance1`

Tunable When disabled

`Oracle_home`
The absolute path of the `ORACLE_HOME` of the web tier component of Oracle Fusion Middleware.

Category Required

Per Node False

Data Type String

Default Null

Tunable When disabled

Attributes

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	`ha-cluster/data-service/opmn`

See Also

`pmfadm(1M)` `scha_resource_get(3HA)`, `clresource(1CL)`, `clresourcetype(1CL)`, `attributes(5)`, `ORCL.ohs(5)`, `r_properties(5)`, `SUNW.gds(5)`

Oracle Solaris Cluster Data Services Planning and Administration Guide, *Oracle Solaris Cluster Data Service for Oracle Web Tier Guide*