Oracle® Solaris Cluster 3.3 Release Notes

3.3 Release Notes

821-1265

September 2010

Oracle Solaris Cluster 3.3 Release Notes

821-1265

Oracle® Solaris Cluster 3.3 Release Notes

Searching Sun Cluster

Index

This area contains Release Notes information for Oracle® Solaris Cluster hardware and software and for Oracle Solaris Cluster Geographic Edition software. The Release Notes are updated when workarounds and other new information becomes available. Please check back often to ensure that you get the latest information.

Select the particular release notes document that you want from the list below:

- (English) Oracle® Solaris Cluster 3.3 Release Notes
- Oracle® Solaris Cluster Geographic Edition 3.3 Release Notes

(English) Oracle® Solaris Cluster 3.3 Release Notes

English

This document is available in the following languages. The English version might have new information that is not yet available in the other versions.

[Español] [Français] [日本語] 简体中文

This document provides the following information for Oracle® Solaris Cluster 3.3 software.

Contents

For information about Oracle Solaris Cluster Geographic Edition 3.3 software, see the Oracle Solaris Cluster Geographic Edition 3.3 Release Notes .

What's New in Oracle Solaris Cluster 3.3 Software

This section provides information related to new features, functionality, and supported products in Oracle Solaris Cluster 3.3 software.

The following new features and functionality are provided in patches to Oracle Solaris Cluster 3.3 software:

- NEW: Support for SWIFT Alliance and Gateway 7.0
- CLI Support to Configure Oracle RAC Instance Proxy Resource Group
- Support for Sun GlassFish Message Queue 4.4

- IP Address is Now Optional for Zone Cluster Nodes
- HA for Oracle Support for Third-Party Volume Managers with Clustered Oracle ASM Disk Groups
- Support for Apache Tomcat 5.5.28
- Oracle Solaris Cluster HA for Oracle TimesTen

The following new features and functionality are provided in the initial Oracle Solaris Cluster 3.3 release:

- NAS NFS File Systems Now Supported in a Zone Cluster
- Failover-File-System Monitoring
- Denying Cluster Services for a Non-Global Zone
- Trusted Extensions With Zone Clusters Support
- Oracle 11g Release 2 Support
- Load-Based Distribution of Resource Groups
- Node Rename
- Power Management User Interface
- Exporting Cluster File Systems to a Zone Cluster Using a Loopback Mount
- Wizard Updates for the `clsetup` Utility and Oracle Solaris Cluster Manager
- Wizard Support for Oracle ASM
- Support for Oracle WebLogic Server 10.3 in Failover and Clustered Configurations
- Enhanced Oracle E-Business Suite 3.2 Agent
- Oracle Solaris Cluster HA for MySQL Cluster
- Oracle Solaris Cluster HA for Oracle Business Intelligence Enterprise Edition
- Oracle Solaris Cluster HA for PeopleSoft Application Server

Top

CLI Support to Configure Oracle RAC Instance Proxy Resource Group

Support is added to the Oracle Solaris Cluster maintenance commands to configure an Oracle Real Application Clusters (Oracle RAC) instance proxy resource group when Oracle ASM is not also configured. This support is delivered starting in the following HA-Oracle patch versions:

- SPARC: 145335-01
- x86: 145336-01

For supporting configuration instructions, see Configuring an Oracle RAC Instance Proxy Resource Group Without ASM (CLI) .

Procedures for removing the Grid Infrastructure resource are also provided. See Removing an Oracle Grid Infrastructure Resource.

Top

Support for Sun GlassFish Message Queue 4.4

The installation restrictions can be removed when CR 6927826 is putback.

The following patches add support for Sun GlassFish Message Queue 4.4 to the HA for Java Message Queue data service:

- SPARC: 145337-02
- x86: 145338-01

After you install the appropriate patch, observe the following installation restrictions for Message Queue 4.4:

- Do not use the GlassFish Message Queue installer.
- Install only the SVR4 packages (not the IPS packages) by using the `pkgadd` command.

Otherwise, you can following procedures as published in Oracle Solaris Cluster Data Service for Sun Java System Message Queue Guide.

Top

IP Address is Now Optional for Zone Cluster Nodes (6961705)

Specifying an IP address and network interface card (NIC) for each zone cluster node is now optional. This new functionality is included in the Oracle Solaris Cluster 3.3 145333-02 Patch, and it is not documented in the Oracle Solaris Cluster 3.3 `clzonecluster`(1CL) man page. If you do not configure an IP address for each zone cluster node, two things will occur:

1. That specific zone cluster will not be able to configure NAS devices for use in the zone cluster. The cluster uses the IP address of the zone cluster node when communicating with the NAS device, so not having an IP address prevents cluster support for fencing NAS devices.
2. The cluster software will activate any Logical Host IP address on any NIC.

HA for Oracle Support for Third-Party Volume Managers with Clustered Oracle ASM Disk Groups

Support is added to the HA for Oracle data service for third-party volume managers configured with clustered Oracle ASM disk groups. This support is delivered in the following HA for Oracle patches:

- SPARC: 145335-01
- x86: 145336-01

For supporting configuration instructions, see Support for Third-Party Volume Managers with Clustered Oracle ASM Disk Groups .

Support for Apache Tomcat 5.5.28

Support is added to the HA for Apache Tomcat data service for Apache Tomcat versions 5.5.28 and greater. This support requires an HA for Apache Tomcat patch and a change in how the data service is configured. See the following information for details:

- `SUNWsctomcat` Resource Probe Fails with Tomcat Version 5.5.28 & 5.5.29 (6964640)
- Data Service for Apache Tomcat Guide

Top

Oracle Solaris Cluster HA for Oracle TimesTen

Oracle Solaris Cluster now supports a data service for Oracle TimesTen In-Memory Database. For information about installing and configuring HA for Oracle TimesTen, see the Oracle Solaris Cluster Data Service for Oracle TimesTen In-Memory Database Guide .

> **Note**
> The following patch is required for this data service:
>
> - SPARC: 145339-01
> - x86: 145340-01

Top

NAS NFS File Systems Now Supported in a Zone Cluster

Network-Attached Storage (NAS) NFS file systems can be used as storage devices for applications that run in a zone cluster with fencing support. You can enable this feature with the Oracle Solaris Cluster Manager GUI or by using the `clnasdevice add -Z` command. For instructions, see Oracle Solaris Cluster Network-Attached Storage Device Manual.

Top

Failover-File-System Monitoring

Oracle Solaris Cluster 3.3 actively monitors mounted file systems and detects whether a mounted file system is truly available to an application, meaning the underlying storage is accessible. When necessary, the cluster performs a corrective failover. See Chapter 2 of the Oracle Solaris Cluster Software Data Services Planning and Administration Guide.

Top

Denying Cluster Services for a Non-Global Zone

You can turn off cluster functionality for a selected non-global zone, so that a root user logged into one of these zones will not be able to discover or disrupt operation of the cluster. For instructions, see How to Deny Cluster Services For a Non-Global Zone

and How to Allow Cluster Services for a Non-Global Zone in the Oracle Solaris Cluster Data Services Planning and Administration Guide.

Top

Trusted Extensions With Zone Clusters Support

Oracle Solaris Cluster on Solaris Trusted Extensions is now supported, but is limited to certain hardware and software configurations. Contact your Oracle representative for the latest information about the currently supported configurations.

For guidelines and configuration procedures, see Oracle Solaris Cluster Software Installation Guide.

Top

Oracle 11g Release 2 Support

Oracle RAC 11g release 2 is supported with Oracle Solaris Cluster. The following restrictions apply to certain Oracle 11g release 2 data-service configurations:

- HA-Oracle and Oracle RAC - The use of Oracle ASM disk groups is supported only on hardware RAID. The use of Oracle ASM disk groups on devices that are managed by a volume manager is not supported.
- HA-Oracle and Oracle RAC - To configure the Oracle ASM resource groups and their resources, use only CLI commands. Do not use Oracle Solaris Cluster Manager or the `clsetup` utility to configure an Oracle ASM resource group.
- Oracle RAC - To configure an Oracle RAC instance proxy resource group, the configuration tools to use are as follows:
  - If the Oracle RAC database does not use Oracle ASM, you must use either the Oracle Solaris Cluster Manager GUI or the `clsetup` utility.
  - If the Oracle RAC database does use Oracle ASM, you must use only the Oracle Solaris Cluster CLI.

Guidelines and procedures in Oracle Solaris Cluster Data Service for Oracle Guide and Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide are valid for Oracle 11g release 2 unless otherwise stated in the documentation.

Top

Load-Based Distribution of Resource Groups

Load-based resource group distribution lets you create a workload distribution policy. This policy considers system capacity and loading while starting or failing over resources. For configuration procedures, see Configuring the Distribution of Resource Group Load Across Nodes in the Oracle Solaris Cluster Software Installation Guide] and How to Configure Load Limits on a Node in the Oracle Solaris Cluster System Administration Guide.

Top

Node Rename

You can use the `clnode rename` command to change the name of a node that is part of an Oracle Solaris Cluster configuration. You must rename the Oracle Solaris hostname before you can rename the node. See How to Rename a Node in the Oracle Solaris Cluster System Administration Guide.

Top

Power Management User Interface

The Oracle Solaris power management user interface can now turn Oracle Solaris Cluster components and systems on and off in order to save energy. See Overview of Administering the Cluster in the Oracle Solaris Cluster System Administration Guide.

Top

Exporting Cluster File Systems to a Zone Cluster Using a Loopback Mount

You can make a cluster file system available to a zone cluster by mounting the file system in the global cluster and then performing a loopback mount in the zone-cluster node. Oracle Solaris Cluster can manage loopback mounts for cluster file systems on UFS and on Veritas File System. See How to Add a Local File System to a Zone Cluster in the Oracle Solaris Cluster

Software Installation Guide] and How to Set Up the HAStoragePlus Resource for Cluster File Systems in the Oracle Solaris Cluster Data Services Planning and Administration Guide.

Top

Wizard Updates for the **`clsetup`** Utility and Oracle Solaris Cluster Manager

The `clsetup` utility and Oracle Solaris Cluster Manager wizards were expanded to support configuring a `SUNW.vucmm_framework` resource group and resources for an Oracle Real Application Clusters configuration. Information about using `SUNW.vucmm_framework`, also called the multiple-owner volume-manager framework, is in Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide.

Top

Wizard Support for Oracle ASM

You can now configure Oracle ASM for HA-Oracle and Oracle RAC by using wizards in the `clsetup` utility and the Oracle Solaris Cluster Manager GUI. You can also configure HA-Oracle and Oracle RAC in a zone cluster. Information is in Oracle Solaris Cluster Data Service for Oracle Guide and Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide.

Top

Support for Oracle WebLogic Server 10.3 in Failover and Clustered Configurations

Oracle Solaris Cluster now supports Oracle WebLogic Server 10.3 in the failover and clustered ("multi-master") modes of operation. This includes support for Oracle WebLogic Server 10.3 servers in zone clusters. Information is located in Oracle Solaris Cluster Data Service for WebLogic Server Guide.

Top

Enhanced Oracle E-Business Suite 3.2 Agent

The Oracle E-Business Suite 3.2 agent now supports Parallel Concurrent Processing. Information is located in Oracle Solaris Cluster Data Service for Oracle E-Business Suite Guide.

Top

Oracle Solaris Cluster HA for MySQL Cluster

Oracle Solaris Cluster now supports a data service for MySQL Cluster. For information about installing and configuring HA for MySQL Cluster, see the Oracle Solaris Cluster Data Service for MySQL Cluster Guide. Additional information to support MySQL Cluster has been added to the existing Oracle Solaris Cluster Data Service for MySQL Guide.

Top

Oracle Solaris Cluster HA for Oracle Business Intelligence Enterprise Edition

Oracle Solaris Cluster now supports a data service for Oracle Business Intelligence Enterprise Edition. For information about installing and configuring HA for Oracle Business Intelligence Enterprise Edition, see the Oracle Solaris Cluster Data Service for Oracle Business Intelligence Enterprise Edition Guide.

Top

Oracle Solaris Cluster HA for PeopleSoft Application Server

Oracle Solaris Cluster now supports a data service for the PeopleSoft application server on SPARC based platforms. For information about installing and configuring HA for PeopleSoft application server, see the Oracle Solaris Cluster Data Service for PeopleSoft Enterprise Guide .

Top

Features Nearing End of Life

There are no features nearing end of life.

Top

Compatibility Issues

This section contains information about Oracle Solaris Cluster compatibility issues with other products.

- Additional Oracle Solaris Cluster framework compatibility issues are documented in "Planning the Oracle Solaris Cluster Configuration" in Oracle Solaris Cluster Software Installation Guide.
- Additional Oracle Solaris Cluster upgrade compatibility issues are documented in "Upgrade Requirements and Software Support Guidelines" in Oracle Solaris Cluster Upgrade Guide.
- For other known problems or restrictions, see Known Issues and Bugs.

Top

Oracle Database 11.2.0.2

For Oracle Database version 11.2.0.2, when using the Oracle Universal Installer (OUI) to install Oracle Grid Infrastructure, you must specify the Oracle ASM discovery string so that only devices that are accessible by the Oracle Grid Infrastructure user will match. If other devices are matched that the Oracle Grid Infrastructure user cannot access, the installation fails.

For example, instead of setting the disk discovery path as `/dev/did/rdsk/d*`, list only the specific disks that are accessible by the Oracle Grid Infrastructure software owner, such as `/dev/did/rdsk/d1s5,/dev/did/rdsk/d2s5`.

For more information, see https://support.us.oracle.com/oip/faces/secure/km/DocumentDisplay.jspx?id=1303209.1.

Top

Undefined Symbol Errors After **`scinstall`** on First Boot on x86

On the Solaris 10 10/09 OS, after running `scinstall` and booting into an x86 cluster for the first time, the node fails to come up and displays multiple "undefined symbol" messages, such as the following:

```
Sep 13 09:42:58 phys-node-1 genunix: [ID 819705 kern.notice]
/kernel/misc/amd64/cl_haci: undefined symbol
Sep 13 09:42:58 phys-node-1 genunix: [ID 826211 kern.notice]  '_memcpy'
Sep 13 09:42:58 phys-node-1 genunix: [ID 472681 kern.notice] WARNING: mod_load: cannot
load module 'cl_haci'
Sep 13 09:42:58 phys-node-1 genunix: [ID 819705 kern.notice]
/kernel/misc/amd64/cl_haci: undefined symbol
...
```

This issue is caused by Solaris bug 6396272, which is fixed in Solaris patch 142910-17. To avoid this problem, follow procedures in "How to Apply a Rebooting Patch (Node)" in the Oracle Solaris Cluster System Administration Guide to apply this patch to each cluster node.

Top

Solaris Upgrade to Solaris 10 9/10 Overwrites the **`hostname.`*`adp`*** Contents (6971566)

When the Solaris OS is upgraded to the Oracle Solaris 10 9/10 release, the upgrade overwrites the `/etc/hostname.`*`adp`* files. To ensure that you do not lose your configuration data, perform the following steps before and after the upgrade:

1. Before you upgrade the OS, make a backup of the contents of each `/etc/hostname.`*`adp`* file.
2. After the upgrade but before you reboot the system, copy the backup contents of each `/etc/hostname.`*`adp`* file to

the current `/etc/hostname.adp` files.

3. After you reboot, if there was IPMP group information in the `/etc/hostname.adp` files, verify that IPMP groups are created correctly. Use either the `cluster` command or `ifconfig` commands.

If you already upgraded the OS to Oracle Solaris 10 9/10 and your `/etc/hostname.adp` files were overwritten, follow procedures in Solaris administration documentation to recreate the files, as well as affected IPMP groups.

Top

Zone Root Path Gets Modified on Doing a Live Upgrade with ZFS Root on **`cluster`** Branded Zone (6852390)

When a cluster that uses a ZFS root file system and which has zone cluster configured is upgraded by using Live Upgrade, the zone root path is modified such that the upgraded zone cannot boot.

To avoid this problem, apply the following patch before you begin live upgrade of the cluster:

- SPARC: 121430-45
- x86: 121431-46

Top

Cluster Zone Won't Boot Up After Live Upgrade on ZFS Root (6955669)

For a global cluster that uses ZFS for the root file system and which has zone clusters configured, when using Live Upgrade to upgrade to Solaris 10 8/10, the upgraded boot environment does not boot.

Contact your Oracle support representative to learn whether a patch or workaround is available.

Top

Webconsole Branding Must Be Changed to Match Oracle Brand (6925641)

The Oracle Solaris Cluster Manager GUI depends on changes to Java Web Console that are not contained in the Solaris 10 11/09 release. The necessary changes are available in the following required patches:

- 125952-20 (SPARC)
- 125953-20 (x86)

Top

The Command **`zoneadmd`** Should Mount **`cluster`** Branded Zones As **`native`** in Alternate Roots: Solaris Upgrade Fails in Presence of Zone Clusters (6874636)

Problem Summary: If a zone cluster is configured on a global-cluster node that runs the Solaris 10 OS, the use of JumpStart to upgrade the OS of that global-cluster node fails with a warning message like the following:

> Warning
> The Solaris Version (Solaris 10) on slice `c1t1d0s0` cannot be upgraded.
>
> A non-global zone could not be mounted.

The problem affects all JumpStart upgrades of the Solaris 10 OS to the Solaris 10 10/09 OS, if there are zone clusters installed on the global-cluster node. For example, an upgrade from the Solaris 10 5/09 OS to the Solaris 10 10/09 OS fails.

> Note
> (1) The problem does not occur if no zone clusters are installed on the cluster being upgraded.
> (2) The problem does not occur if there are zone clusters, but they exist only in the `Configured` state.

Workaround: Add two files to the Solaris installation miniroot image on the JumpStart server, as shown in the steps below. In

these instructions, server is the name of the JumpStart server that you use to install the upgrade. Perform all steps as superuser.

Perform Step 1 from one node of the cluster that is to be upgraded.

1. Copy the zone-cluster `config.xml` and `platform.xml` files from a node that is to be upgraded to a location on the JumpStart server.

```
# cp /usr/lib/brand/cluster/config.xml.upgrade /net/server/some_dir/config.xml
# cp /usr/lib/brand/cluster/platform.xml /net/server/some_dir/platform.xml
```

Perform Steps 2-6 on the JumpStart server.

2. Create a new Solaris installation image.

```
# cd path_to_Solaris_install_image/Solaris_10/Tools
# ./setup_install_server path_to_new_Solaris_install_image
# cd
```

3. Unpack the miniroot.

```
# /boot/solaris/bin/root_archive unpackmedia path_to_new_Solaris_install_image
path_to_miniroot
```

4. Put the zone-cluster brand files in the miniroot.

```
# mkdir path_to_miniroot/usr/lib/brand/cluster
# cp /net/server/some_dir/config.xml path_to_miniroot/usr/lib/brand/cluster
# cp /net/server/some_dir/platform.xml path_to_miniroot/usr/lib/brand/cluster
```

5. Pack the miniroot and place it in the new installation image.

```
# /boot/solaris/bin/root_archive packmedia path_to_new_Solaris_install_image
path_to_miniroot
```

6. Run the `addclient` command from the new installation image location for each cluster node to be upgraded.

```
# cd path_to_new_Solaris_install_image/Solaris_10/Tools
# ./addclient your_configuration_details
```

Perform this workaround just once for each JumpStart server and for each Solaris update to which you plan to upgrade.

Top

Accessibility Features for People With Disabilities

To obtain accessibility features that have been released since the publishing of this media, consult Section 508 product assessments that are available from Oracle on request to determine which versions are best suited for deploying accessible solutions.

Top

Solaris Volume Manager GUI

The Enhanced Storage module of Solaris Management Console (Solaris Volume Manager) is not compatible with Oracle Solaris Cluster software. Use the command-line interface or Oracle Solaris Cluster utilities to configure Solaris Volume Manager software.

Top

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.

Top

Product Name Changes

This section provides information about product name changes for applications that Oracle Solaris Cluster software supports. Depending on the Oracle Solaris Cluster software release that you are running, your Oracle Solaris Cluster documentation might not reflect the following product name changes.

Top

Current Product Name	Former Product Name
Oracle Solaris Cluster	Sun Cluster (also Solaris Cluster and Java Availability Suite)
Oracle Solaris Cluster Geographic Edition	Sun Cluster Geographic Edition
Oracle Solaris Cluster data services	Sun Cluster data services
Oracle Solaris Cluster Data Service for Oracle Grid Engine	Sun Cluster Data Service for Sun Grid Engine
Oracle Solaris Cluster Data Service for Oracle VM Server for SPARC	Sun Cluster Data Service for LDoms Guest Domain
Oracle Solaris Cluster Manager	Sun Cluster Manager

Top

Supported Products

This section describes the supported software and memory requirements for Oracle Solaris Cluster 3.3 software.

Top

Data Services

Contact your Oracle sales representative for the complete list of supported data services (agents) and application versions. Links to the documentation for many data service agents are available here.

Top

File Systems

Solaris 10 SPARC

File System	Additional Information

Solaris UFS	
Solaris ZFS	Not supported for the `/globaldevices` file system
Sun StorEdge QFS	
Sun QFS 5.0 and 5.1 - Standalone file system	Supported Data Services: All failover data services External Volume Management: Solaris Volume Manager only
QFS 4.6 - Standalone file system	Supported Data Services: All failover data services External Volume Management: Solaris Volume Manager, VxVM
QFS 4.6, 5.0, and 5.1 - Shared QFS file system	Supported Data Services: Oracle RAC External Volume Management: Solaris Volume Manager for Sun Cluster
QFS 4.6, 5.0, and 5.1 - Shared QFS clients outside the cluster (SC-COTC)	Supported Data Services: None; only a shared file system is supported External Volume Management: No external volume manager is supported
QFS 4.6, 5.0, and 5.1 - HA-SAM Failover	Supported Data Services None; only a shared file system is supported External Volume Management: No external volume manager is supported
Veritas File System components that are delivered as part of Veritas Storage Foundation 5.0 and 5.1. Version 5.0 requires a minimum of MP3 RP3.	

Top

Solaris 10 x86

File System	Additional Information
Solaris UFS	
Solaris ZFS	Not supported for the `/globaldevices` file system
Sun StorEdge QFS	
Sun QFS 5.0 and 5.1 - Standalone file system	Supported Data Services: All failover data services External Volume Management: Solaris Volume Manager only
QFS 4.6 - Standalone file system	Supported Data Services: All failover data services External Volume Management: Solaris Volume Manager, VxVM
QFS 4.6, 5.0, and 5.1 - Shared QFS file system	Supported Data Services: Oracle RAC External Volume Management: Solaris Volume Manager for Sun Cluster
QFS 4.6, 5.0, and 5.1 - Shared QFS clients outside the cluster (SC-COTC)	Supported Data Services: None; only a shared file system is supported External Volume Management: No external volume manager is supported

QFS 4.6, 5.0, and 5.1 - HA-SAM Failover	Supported Data Services None; only a shared file system is supported *External Volume Management:* No external volume manager is supported
Veritas File System components that are delivered as part of Veritas Storage Foundation 5.0 and 5.1. Version 5.0 requires a minimum of MP3 RP3.	

Top

Memory Requirements

Oracle Solaris Cluster 3.3 software requires the following memory requirements for every cluster node:

- Minimum of 1 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.

Top

Solaris Operating System (OS)

Oracle Solaris Cluster 3.3 software and Quorum Server software requires one of the following versions of the Solaris OS:

- Solaris 10 - Solaris 10 10/09, Oracle Solaris 10 9/10

> Note -
> Oracle Solaris Cluster 3.3 software does not support multiple versions of Solaris software in the same running cluster.

Top

Sun Logical Domains (LDoms)

This Oracle Solaris Cluster release supports Sun Logical Domains 1.2 and 1.3 software. For version 1.2, a minimum of patch 142840-05 is required.

Top

Sun Management Center

This Oracle Solaris Cluster release supports Sun Management Center software versions 3.6.1 and 4.0.

Top

Sun StorageTek Availability Suite

This Oracle Solaris Cluster release supports Sun StorageTek Availability Suite 4.0 software. Support requires a minimum of patch 123246-05 for SPARC or 123247-05 for x86.

Top

Volume Managers

This Oracle Solaris Cluster release supports the following volume managers.

Solaris 10 SPARC

Volume Manager	Cluster Feature
Solaris Volume Manager	Solaris Volume Manager for Sun Cluster
Veritas Volume Manager (VxVM) components that are delivered as part of Veritas Storage Foundation 5.0 and 5.1. Version 5.0 requires a minimum of MP3 RP3.	VxVM 5.0 and VxVM 5.1 cluster feature (with RAC only)

Top

Solaris 10 x86

Volume Manager	Cluster Feature
Solaris Volume Manager	Solaris Volume Manager for Sun Cluster
Veritas Volume Manager components that are delivered as part of Veritas Storage Foundation 5.0 and 5.1. Version 5.0 requires a minimum of MP3 RP3. Veritas Volume Replicator (VVR) and Fast Mirror Resynchronization (FMR) are not supported on the x86 platform.	Not applicable - Oracle Solaris Cluster 3.3 software does not support the VxVM cluster feature on the x86 platform.

Top

Oracle Solaris Cluster Security Hardening

Oracle Solaris Cluster Security Hardening uses the Solaris Operating System hardening techniques recommended by the Sun BluePrints™ program to achieve basic security hardening for clusters. The Solaris Security Toolkit automates the implementation of Oracle Solaris Cluster Security Hardening.

The Oracle Solaris Cluster Security Hardening documentation is available http://blogs.sun.com/security/entry/reference_security_blueprints . From this URL, scroll down to the 2003 heading to locate the article "Securing the Sun Cluster 3.x Software." The documentation describes how to secure Sun Cluster 3.x deployments in a Solaris environment. The description includes the use of the Solaris Security Toolkit and other best-practice security techniques recommended by Oracle security experts. The following data services are supported by Oracle Solaris Cluster Security Hardening:

- Oracle Solaris Cluster HA for Apache
- Oracle Solaris Cluster HA for Apache Tomcat
- Oracle Solaris Cluster HA for DHCP
- Oracle Solaris Cluster HA for DNS
- Oracle Solaris Cluster HA for MySQL
- Oracle Solaris Cluster HA for NFS
- Oracle Solaris Cluster HA for Oracle
- Oracle Solaris Cluster HA for Oracle E-Business Suite
- Oracle Solaris Cluster HA for Oracle Grid Engine
- Oracle Solaris Cluster Support for Oracle Real Application Clusters
- Oracle Solaris Cluster HA for PostgreSQL
- Oracle Solaris Cluster HA for Samba
- Oracle Solaris Cluster HA for Siebel
- Oracle Solaris Cluster HA for Solaris Containers
- Oracle Solaris Cluster HA for SWIFTAlliance Access
- Oracle Solaris Cluster HA for SWIFTAlliance Gateway
- Oracle Solaris Cluster HA for Sun Java System Directory Server
- Oracle Solaris Cluster HA for Sun Java System Message Queue
- Oracle Solaris Cluster HA for Sun Java System Messaging Server
- Oracle Solaris Cluster HA for Sun Java System Web Server
- Oracle Solaris Cluster HA for Sybase ASE
- Oracle Solaris Cluster HA for WebLogic Server
- Oracle Solaris Cluster HA for WebSphere MQ
- Oracle Solaris Cluster HA for WebSphere MQ Integrator

Top

Known Issues and Bugs

The following known issues and bugs affect the operation of the Oracle Solaris Cluster 3.3 release. Bugs and issues are grouped into the following categories:

- Administration
- Data Services
- GUI
- Installation
- Localization
- Upgrade

Also see Compatibility Issues and Documentation Issues .

Top

Administration

Oracle's SPARC T3-4 Fails During Reboot (6993321)

Problem Summary: During a reboot, Oracle's SPARC T3-4 server with four processors fails to connect to the Oracle Solaris Cluster framework. For more detailed information, see 6986241. Error messages similar to the following appear:

```
Sep 20 15:18:53 svc.startd [8]: svc:/system/pools:default: Method or service exit timed out. Killing contract 29.
Sep 20 15:18:53 svc.startd [8]: svc:/system/pools:default: Method "/lib/svc/method/svc-pools start" failed due to
signal KILL.
...
Sep 20 15:20:55 solta svc.startd [8]: system/pools:default failed: transitioned to maintenance (see 'svcs -xv' for
details)
...
Sep 20 15:22:12 solta INITGCHB: Given up waiting for rgmd.
...
Sep 20 15:23:12 solta Cluster.GCHB_resd: GCHB system error: scha_cluster_open failed with 18
Sep 20 15:23:12 solta : No such process
```

Workaround: Use the `svccfg` command to increase the service timeout to 300 seconds. Boot into non-cluster mode and perform the following commands:

```
svccfg -s svc:/system/pools setprop start/timeout_seconds = 300
svcadm refresh svc:/system/pools
```

After you perform these commands, boot into cluster mode.

Top

Default Quorum Timeout of 25 Seconds May Not Be Adequate for Remote Quorums on IP Networks With Routing Involved (6974930)

Problem Summary: Cluster reconfiguration has a timeout of 25 seconds for the completion of quorum operations. This timeout value is deemed appropriate for typical quorum-device connectivity technologies, such as Fiber Channel or Direct-Attached SCSI. But a 25-second timeout might not be adequate for other technologies that are based on IP networks, especially with routing in such networks. This caveat also applies to any SCSI device, regardless of topology, that might be experiencing I/O timeouts and retrys, even if the latter eventually do succeed.

If you see failures of cluster nodes with the message `CMM: Unable to acquire the quorum device`, the cluster reconfiguration's quorum operations did not complete within the default timeout of 25 seconds. This might be a problem with the quorum device or with the path to the device. Check that both the quorum device and the path to it are functional.

Workaround: If the problem continues and you want to keep using the same quorum configuration, perform the following steps to increase the quorum operations timeout on all nodes of the cluster. The example below shows how to increase the timeout.

1. Become superuser.
2. On each cluster node, edit the `/etc/system` file as superuser to set the timeout to a high value.
   The following example sets the timeout to 700 seconds.

   ```
   phys-schost# vi /etc/system
   ...
   set cl_haci:qd_acquisition_timer=700
   ```

3. From one node, shut down the cluster.

   ```
   phys-schost-1# cluster shutdown -g0 -y
   ```

4. Boot each node back into the cluster.
   Changes to the `/etc/system` file are initialized after the reboot.

NOTE
For Oracle Real Application Clusters (Oracle RAC), do not change the default quorum timeout of 25 seconds. In certain split-brain scenarios, a longer timeout period might lead to the failure of Oracle RAC VIP failover, due to the VIP resource timing out. If the quorum device being used is not conforming with the default 25 seconds timeout, use a different quorum device.

Top

The **`cluster status -t node`** Command Does Not Work (6973987)

Problem Summary: The `cluster status -t node` command does not work as expected and gives the following error message:

```
% cluster status -t node
cluster: (C103603) Object type "node" does not have status.
```

Workaround: Run the following command instead to query the comprehensive node status:

```
% clnode status [-v] [-l] [-m] [-Z zonecluster | global | all ] [+ | node ...]
```

Top

Newly Added Zone-Cluster Node Is Not Recognized by **`clresourcegroup`** or by Some **`clzonecluster`** Options (6957622)

Problem Summary: This issue occurs for a zone cluster that is configured on a subset of nodes in the global cluster and that has at least one zone-cluster node in the Online state. If a new zone-cluster node is added to that zone cluster, operations related to resource groups and resources do not recognize the newly added node as a valid node of the zone cluster. For example, an attempt to create a resource group with a node list that contains the newly added node results in an `Invalid node` error. Similarly, if you run the command `scha_cluster_get -O ALL_ZONES -Z` *`zonecluster`* from the global zone, the output does not report the newly added zone-cluster node.

Workaround: After a new node is added to a zone cluster, you must reboot all global-cluster nodes before they will recognize the new zone-cluster node.

To avoid this issue, ensure that all zone-cluster nodes that you might need are created at the time you create the zone cluster itself.

Top

When Using ZFS on Replicated Devices, Device Group Can Be Moved Away, Leaving Application Unable to Write to Storage (6848624)

Problem Summary: In a campus cluster configuration using Hitachi Universal Replicator and ZFS, if a user issues a manual switchover of a device group containing replicated ZFS storage pools (zpools), the operation results in the following undesirable state:

- The zpools do not migrate to the new primary node and will have their state changed to write-disabled.
- The application resource groups using the zpools do not migrate to the new primary node.
- The associated raw-disk device groups migrate to the new primary node and have their state set to read-write access.

Workaround: Do not manually switch over a device group that contains zpools that are replicated by Hitachi Universal Replicator.

Solaris Cluster Manager Fails to Come Up in a 16-Node Cluster (6594485)

Problem Summary: For a 16-node cluster, the Oracle Solaris Cluster Manager GUI is not usable.

Workaround: Use instead the `clsetup` utility or the Oracle Solaris Cluster administrative commands.

Top

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.

Top

Load Balancing Weight Is Not Updated if the Scalable Resource Is Disabled (6883814)

Problem Summary: If the load balancing weight of a scalable resource is changed when the resource is in a disabled state, the changed weight is not reflected in the client request distribution after the resource is enabled.

Workaround: Execute the command again after the resource is enabled.

Top

Deleting a Resource Which Has Dependent **`clrs: (C979102) "<rs>" is not present in the property list`** (6970087)

Public Summary: Unable to delete a resource which has dependencies that were set by using the `-F` option in the `clresourcegroup delete` command.

Workaround: Unset the dependencies before you delete the resources.

Top

The **`global_fencing`** Property Code is Broken When the Value is Changed to **`prefer3`** (6879360)

Problem Summary: Oracle Solaris Cluster attempts to verify that a storage device fully supports SCSI-3 PGR before allowing the user to set its fencing property to `prefer3`. This verification might succeed when it should fail.

Workaround: Ensure that a storage device is certified by Oracle Solaris Cluster for use with SCSI-3 PGR before changing the fencing setting to `prefer3`.

Top

Global-Cluster Node Cannot Get Zone Cluster Resource Group or Resource Information if its Zone-Cluster Node Has Never Been Booted Into Cluster Mode (6770391)

Problem Summary: This problem occurs for a zone cluster that is configured with resources and resource groups. If a zone-cluster node has not booted into cluster mode even once, then running the `clresourcegroup status -Z` *`zonecluster`* and `clresource status -Z` *`zonecluster`* commands in the global zone returns an `unexpected error` message. Similarly, running the `scha_resourcegroup_get -Z` *`zonecluster`* and `scha_resource_get -Z` *`zonecluster`* commands in the global zone returns error 18, if the zone-cluster nodes are not yet booted into cluster mode. On such nodes, Oracle Solaris Cluster Manager does not report the status of resource groups and resources for the zone cluster.

Workaround: Boot the zone-cluster nodes into cluster mode by running the following command in the global zone, where *`nodelist`* is the comma-separated list of nodes on which the zone-cluster nodes are not yet booted into cluster mode:

```
# clzonecluster boot -n nodelist zonecluster
```

This command starts the cluster services that are required to work around this problem. You only need to boot the zone-cluster nodes once when the physical node is booted. After that, you can halt or boot the zone-cluster nodes at your discretion, for as long as the physical node remains up.

Top

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.

Top

More Validation Checks Needed When Combining DIDs (6605101)

Problem Summary: The `scdidadm` and `cldevice` commands are unable to verify that replicated SRDF devices 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 DID device instances specified are replicas of each other and that they belong to the specified replication group.

Top

Data Services

HAStoragePlus Resource Configured for Zone Cluster with VxFS File System Fails to Come Online (6972831)

Problem Summary: When a VxFS file system is configured with an HAStoragePlus resource for a zone cluster, the resource fails to come online because of a file-system check failure. The error message is similar to the following:

```
File system check of mountpoint (rawdevice) failed: (39)
rawdevice: BADSUPERBLOCK AT BLOCK 16: MAGIC NUMBER WRONG
rawdevice: USE AN ALTERNATE SUPERBLOCK TO SUPPLY NEEDED INFORMATION;
rawdevice: e.g. fsck [-F ufs] -o b=# [special ...]
rawdevice: where # is the alternate super block. SEE fsck_ufs(1M).
rawdevice: UNEXPECTED INCONSISTENCY; RUN fsck MANUALLY.
Failed to fsck : mountpoint
```

Workaround: Ensure that the VxFS file system information is kept in the `/etc/vfstab` file of the global zone along and that the file system is added to the zone cluster configuration by using the `clzonecluster` command.

Top

11gR2 CRS May Not Start if **`crs_framework`** Resource Was Disabled During Boot Up (6971060)

Problem Summary: This issue applies to cluster nodes that have Oracle 11g release 2 installed and are using the RAC framework.

If a node is rebooted with the RAC framework disabled, the node upon booting attempts to start the Oracle Grid Infrastructure. This results in the cluster node being in a state where the Grid Infrastructure cannot be started or stopped until the node is rebooted, either with the RAC framework enabled or the Oracle Grid Infrastructure autostart disabled.

Workaround: Before you reboot a node with the RAC framework disabled, first disable the Oracle Grid Infrastructure autostart. See the Oracle 11g release 2 documentation for the command to use to disable autostart of the Grid Infrastructure.

Top

The Oracle WebLogic Server Proxy Server Resource Fails to Start (6970737)

Problem Summary: The Oracle WebLogic Server Proxy Server resource fails to start if the Managed Servers are offline.

Workaround: The WebLogic Server Proxy Server resource within a WebLogic Server proxy resource group should be dependent on WebLogic Server Managed Server resources. Strong resource dependencies ensure that the WebLogic Server Proxy Server is brought online before Managed Server resources are brought online. Use this command to specify the dependencies between the WebLogic Server proxy resource and the Managed Server resource:

```
# clresource create \
-t SUNW.wls \
-g wls-proxy-rg \
-p resource_dependencies=wls-mgd1-rs,wls-mgd2-rs ... \
wls-proxy-rs

# clresource set \
-p resource_dependencies=wls-mgd1-rs,wls-mgd2-rs ... \
wls-proxy-rs
```

Top

The Oracle WebLogic Server Stops Core Dumps (6970466)

Problem Summary: The Oracle WebLogic Server stops core dumps if the `smooth_shutdown` property is set to `true` and the `server_url` property does not include the port number.

Workaround: The WLS `server_url` property value should be in the form of `http://`*`hostname`*`:`*`port`*.

_

Top

Cannot Create RAC Proxy Resource When User Has the Same Group ID for Grid and Database Homes (6969233)

Problem Summary: The `SUNW.scalable_rac_server_proxy` resource cannot be created if the software installations of both the Oracle 11g release 2 Database and Grid Infrastructure used the same DBA group ID.

Workaround: Reinstall the Oracle 11g release 2 Database software using a different DBA group ID than that which was used for the Grid Infrastructure installation.

Top

The **`SUNW.crs_framework`** Resource Does Not Attempt a Restart When Oracle Clusterware Goes Offline (6966668)

Problem Summary: If Oracle Clusterware software fails on a node, the `SUNW.crs_framework` resource type does not automatically restart the software. This problem occurs whenever Oracle Clusterware software itself goes down and is unable to restart itself.

Workaround: Restart the Oracle Clusterware software manually. See the appropriate Oracle release documentation to determine the appropriate commands.

Top

`SUNWsctomcat` Resource Probe Fails with Tomcat Version 5.5.28 & 5.5.29 (6964640)

Problem Summary: Starting with version 5.5.28, the Apache Tomcat server changed the behavior of the `http get` command.

Workaround: Due to this change, an alternative probe algorithm is supplied in the following Oracle Solaris Cluster patches for HA-Apache Tomcat:

- SPARC: 145341-01
- x86: 145342-01

This probe algorithm uses the `wget` command, bypassing the normal probe. An HA for Apache Tomcat data service using Apache Tomcat version 5.5.28 or greater must use the `wget` probe algorithm. You can use the `wget` probe algorithm for versions earlier than 5.5.28 as well.

In addition, supporting documentation for this functionality is added to or replaces what is published in Oracle Solaris Cluster Data Service for Apache Tomcat Guide. See Data Service for Apache Tomcat Guide for details.

Top

HAStoragePlus Resource Configured in Scalable Resource Group with Cluster File System Stays at "Starting" State Indefinitely (6960386)

Problem Summary: When the `/etc/vfstab` file entry for a cluster file system has a mount-at-boot value of "no" and the cluster file system is configured in a `SUNW.HAStoragePlus` resource that belongs to a scalable resource group, the `SUNW.HAStoragePlus` resource fails to come online. The resource stays in the Starting state until `prenet_start_method` is timed out.

Workaround: In the `/etc/vfstab` file's entry for the cluster file system, set the mount-at-boot value to "yes".

Top

Configuration of Resource Groups With `start`*`app_name`* Script Fails (6941251)

Problem Summary: The `start`*`app_name`* script fails to create resource groups, because an incorrect `nodelist` is passed to the `clresourcegroup create` command.

Workaround: You must specify `-n nodelist`, even for the default node list.

Top

Scalable Resource Fails To Start Due To Failure of Unrelated IPMP Group (6938555)

Problem Summary: A scalable resource that depends on a `SUNW.SharedAddress` resource fails to come online due to a failure of an IPMP group that is on a subnet not used by the shared-address resource. The following messages can be 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 scalable resource that failed to start before.

Alternatively, remove the failed IPMP group (see the `ifconfig`(1M) man page) and restart the scalable resource that failed to start before.

> **NOTE**
> Network interfaces in the failed IPMP group will no longer be monitored by Oracle Solaris Cluster after the IPMP group is removed. Resources that depend on the failed IPMP group can no longer be created or started until the IPMP group is recreated and repaired.

Top

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 running in different zone clusters.

Workaround: The scalable applications should not be configured to bind to `INADDR_ANY` as the local IP address or they should bind to a port that does not conflict with another scalable application.

Top

Validation Should Fail for a Scalable Resource With the **`Outgoing_connection`** Property Set and With Multiple Entry Points (6886818)

Problem Summary: The problem occurs when the `Outgoing_Connection` property is set on a scalable resource and the resource is configured with multiple failover resource groups, each containing one or more scalable addresses. Such a configuration is not caught during validation, but TCP connections for the scalable service might be disrupted as a result.

Workaround: If the `Outgoing_Connection` property is enabled for a scalable resource, include all its scalable addresses in a single failover resource group.

Top

GUI

Data Service Configuration Wizards for Apache, NFS, HA-Oracle, and SAP Do Not Set an Offline-Restart Dependency on an HAStoragePlus Resource (6947273)

Problem Summary: The data-service configuration wizards in the Oracle Solaris Cluster GUI or the `clsetup` utility do not set an offline-restart dependency on an HAStoragePlus resource from the application. This is a regression in the wizards that was introduced in this release with the new failover-file-system monitoring feature.

If the resource is an existing resource that was previously created by the wizard, it will continue to operate without failover-file-system monitoring until the user explicitly upgrades the HAStoragePlus resource type.

Workaround: To enable failover-file-system monitoring for an existing resource, , take offline the HAStoragePlus resource, upgrade it to the 3.3 version resource type, and modify its resource dependencies appropriately.

To create a new resource with failover-file-system monitoring enabled, do not use the data-service wizards to create the HAStoragePlus resources where file-system monitoring is to be enabled. Use the command-line interface instead.

Top

RAC Wizard Configures QFS MDS Resources for Different Zone Clusters Under Same QFS MDS Resource Group by Default (6887179)

Problem Summary: The RAC wizard, by default, creates the QFS MDS resources for different zone clusters and the global cluster under the same QFS MDS resource group. The zone cluster nodes can be halted or rebooted independently of each other. So if the resources are in the same resource group, the administration of one resource could inadvertently impact another.

Workaround: The wizard provides an option to edit the auto-generated resource and resource group names in the review panel. Use this option to edit the QFS MDS resource group name, which is auto-generated by the wizard, and give it a value that is different from the one that already exists on the cluster.

Top

Installation

Autodiscovery Does Not Work for the **`qlge`** Driver for PCIe FCoE CNA (6939847)

Problem Summary: During an Oracle Solaris Cluster installation, auto discovery for the `qlge` driver for the PCIe LP and PCIe ExpressModule FCoE Converged Network Adapters (CNAs) does not work. The following products are affected:

- Oracle's Sun Storage 10GbE FCoE PCIe CNA
- Oracle's Sun Storage 10GbE FCoE ExpressModule CNA

Workaround: When you run the `scinstall` utility and you are prompted for the interconnect adapters, select Other and type the name of each `qlge` interface.

Top

Localization

Result of System Requirements Checking Is Wrong (6495984)

Problem Summary: When you use the `installer` utility in the Simplified Chinese and Traditional Chinese locales to install Oracle Solaris Cluster software, the software that checks the system requirements incorrectly reports that the swap space is 0 Mbytes.

Workaround: Ignore this reported information. In these locales, you can run the following command to determine the correct swap space:

```
# df -h | grep swap
```

Top

Upgrade

The Command **`zoneadmd`** Should Mount **`cluster`** Branded Zones As **`native`** in Alternate Roots: Solaris Upgrade Fails in Presence of Zone Clusters (6874636)

Problem Summary: This problem affects all JumpStart upgrades of the Oracle Solaris 10 OS to at least the Solaris 10 10/08 OS, if there are zone clusters installed on the global-cluster node. If a zone cluster is configured on a global-cluster node that runs the Oracle Solaris 10 OS, the use of JumpStart to upgrade the OS of that global-cluster node fails with a warning message similar to the following:

```
Warning
The Oracle Solaris Version (Solaris 10) on slice c1t1d0s0 cannot be upgraded.

A non-global zone could not be mounted.
```

For example, an upgrade from the Solaris 10 5/09 OS to the Solaris 10 10/09 OS fails.

Note:
(1) The problem does not occur if no zone clusters are installed on the cluster being upgraded.
(2) The problem does not occur if there are zone clusters, but they exist only in the `Configured` state.

Workaround: Add two files to the Oracle Solaris installation miniroot image on the JumpStart server, as shown in the following steps. In these instructions, server is the name of the JumpStart server that you use to install the upgrade. Perform all steps as superuser.

Perform Step 1 from one node of the cluster that is to be upgraded.

1. Copy the zone-cluster `config.xml` and `platform.xml` files from a node that is to be upgraded to a location on the JumpStart server.

```
# cp /usr/lib/brand/cluster/config.xml.upgrade /net/server/some_dir/config.xml
# cp /usr/lib/brand/cluster/platform.xml /net/server/some_dir/platform.xml
```

Perform Steps 2-6 on the JumpStart server.

2. Create a new Oracle Solaris installation image.

```
# cd path_to_Solaris_install_image/Solaris_10/Tools
# ./setup_install_server path_to_new_Solaris_install_image
# cd
```

3. Unpack the miniroot.

```
# /boot/solaris/bin/root_archive unpackmedia path_to_new_Solaris_install_image
path_to_miniroot
```

4. Put the zone-cluster brand files in the miniroot.

```
# mkdir path_to_miniroot/usr/lib/brand/cluster
# cp /net/server/some_dir/config.xml path_to_miniroot/usr/lib/brand/cluster
# cp /net/server/some_dir/platform.xml path_to_miniroot/usr/lib/brand/cluster
```

5. Pack the miniroot and place it in the new installation image.

```
# /boot/solaris/bin/root_archive packmedia path_to_new_Solaris_install_image
path_to_miniroot
```

6. Run the `addclient` command from the new installation image location for each cluster node to be upgraded.

```
# cd path_to_new_Solaris_install_image/Solaris_10/Tools
# ./addclient your_configuration_details
```

Perform this workaround just once for each JumpStart server and for each Oracle Solaris update to which you plan to upgrade.

Top

Setting the `num_zoneclusters` Property Causes a Core Dump on a Cluster That Was Upgraded from Sun Cluster 3.1 to Sun Cluster 3.2 1/09 Software (6846142)

Problem Summary: The property `num_zoneclusters` that was introduced in Sun Cluster 3.2 1/09 software defines the maximum number of zone clusters that can be created on the cluster. When a cluster is upgraded to Oracle Solaris Cluster 3.3 software from a version earlier than the Sun Cluster 3.2 1/09 release, you must set this property by using the `cluster` command before you configure any zone clusters. However, if the cluster is upgraded to Oracle Solaris Cluster 3.3 from Sun Cluster 3.1 software, the process of setting the `num_zoneclusters` property generates a core dump.

Workaround: Set the properties `max_nodes` and `max_privatenets` before you set the `num_zoneclusters` property. Use the following procedure:

1) Reboot all nodes into non-cluster mode.
2) Use the `/usr/cluster/bin/clsetup` or `/usr/cluster/bin/cluster` commands to set the values of the properties `max_nodes` and `max_privatenets`.
3) Reboot all nodes into cluster mode.
4) Run the `/usr/cluster/bin/cluster` command to set the value of `num_zoneclusters`.
Use the following command-line syntax, where `172.16.0.0` is an example private-network address:

```
# cluster set-netprops \
-p private_netaddr="172.16.0.0" \
-p max_nodes=number_of_nodes_expected_in_the_cluster \
-p max_privatenets=number_of_networks_expected_in_the_cluster
```

5) Start the `clsetup` utility and make selections similar to those shown:

```
# clsetup
*** Main Menu ***


1) Change Network Addressing and Ranges for the Cluster Transport
2) Show Network Addressing and Ranges for the Cluster Transport

?) Help with menu options
q) Quit

    Option: 1

>>> Change Network Addressing and Ranges for the Cluster Transport <<<


Network addressing for the cluster transport is currently configured
as follows:

Private Network ===
 private_netaddr: 172.16.0.0
   private_netmask: 255.255.240.0
   max_nodes: 64
   max_privatenets: 10
   num_zoneclusters: 12


Do you want to change this configuration (yes/no) [yes]? yes

The default network address for the cluster transport is 172.16.0.0.

Do you want to use the default (yes/no) [yes]? yes

The default netmask for the cluster transport is 255.255.240.0.

Do you want to use the default (yes/no) [yes]? no

The combination of private netmask and network address will dictate
both the maximum number of nodes and private networks that can be
supported by a cluster. Given your private network address, this
program will generate a range of recommended private netmasks based on
the maximum number of nodes and private networks that you anticipate
for this cluster.

In specifying the anticipated number of maximum nodes and private
networks for this cluster, it is important that you give serious
consideration to future growth potential. While both the private
netmask and network address can be changed later, the tools for making
such changes require that all nodes in the cluster be booted into
noncluster mode.

Maximum number of nodes anticipated for future growth [64]? 64

Maximum number of private networks anticipated for future growth [10]? 10

Specify a netmask of 255.255.248.0 to meet anticipated future
requirements of 64 cluster nodes and 10 private networks.

To accommodate more growth, specify a netmask of 255.255.240.0 to
support up to 64 cluster nodes and 20 private networks.

What netmask do you want to use [255.255.248.0]? 255.255.248.0
Is it okay to proceed with the update (yes/no) [yes]? yes
```

Top

Patches and Required Firmware Levels

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

- Applying an Oracle Solaris Cluster 3.3 Core Patch
- Removing an Oracle Solaris Cluster 3.3 Core Patch
- Patch Management Tools
- Patch for Cluster Support for Sun StorageTek 2530 Array
- My Oracle Support

If you are upgrading to Oracle Solaris Cluster 3.3 software, see the Oracle Solaris Cluster Upgrade Guide. Applying an Oracle Solaris Cluster 3.3 Core patch does not provide the same result as upgrading the software to the Oracle Solaris Cluster 3.3 release.

> **Note**
> Read the patch README file before applying or removing any patch.

You must be a registered My Oracle Support user to view and download the required patches 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 My Oracle Support .

Top

Applying an Oracle Solaris Cluster 3.3 Core Patch

Complete the following procedure to apply an Oracle Solaris Cluster 3.3 core patch. Ensure that all nodes of the cluster are maintained at the same patch level.

How to Apply an Oracle Solaris Cluster 3.3 Core Patch

> **Caution**
> If an Oracle Solaris Cluster 3.3 core patch is removed, any resources that were upgraded in Step 3 must be downgraded to the earlier resource type versions. The procedure for downgrading requires planned downtime of these services. Therefore, do not perform Step 3 until you are ready to commit the Oracle Solaris Cluster 3.3 core patch permanently to your cluster.

1. Install the patch by using the usual rebooting patch procedure for a core patch.
2. Verify that the patch has been installed correctly on all nodes and is functioning properly.
3. Perform Resource Type upgrades to any new versions of the resource type available. Run `clsetup` to obtain the list of new resource types. For more details, see the documentation on "Upgrading a Resource Type" in the Oracle Solaris Cluster Data Services Planning and Administration Guide.

For information about registering a resource type, see "Registering a Resource Type" in the Oracle Solaris Cluster Data Services Planning and Administration Guide.

Top

Removing an Oracle Solaris Cluster 3.3 Core Patch

Complete the following procedure to remove an Oracle Solaris Cluster 3.3 core patch.

How to Remove an Oracle Solaris Cluster 3.3 Core Patch

1. List the resource types on the cluster.

```
# clresourcetype list
```

2. If you have upgraded to any new resource types after applying the core patch, follow the directions in "How to Remove a Resource Type" in the Oracle Solaris Cluster Data Services Planning and Administration Guide.
3. Remove the Oracle Solaris Cluster 3.3 core patch from each node on which you installed the patch.

```
# patchrm patch-id
```

4. Reboot into cluster mode all of the nodes from which you removed the Oracle Solaris Cluster 3.3 core patch. Rebooting all of the nodes from which you removed the Oracle Solaris Cluster 3.3 core patch before rebooting any unaffected nodes ensures that the cluster is formed with the correct configuration information on all nodes. If all nodes on the cluster were patched with the core patch, you can reboot the nodes into cluster mode in any order.
5. Reboot any remaining nodes into cluster mode.

For instructions about rebooting nodes into cluster mode, see "How to Reboot a Cluster Node" in the Oracle Solaris Cluster System Administration Guide.

Top

Patch Management Tools

Information about patch management options for the Solaris OS is available at the web site for Oracle Enterprise Manager Ops Center (formerly Sun Ops Center).

The following tools are part of the Solaris OS. Refer to the version of the manual that is published for the Solaris OS release that is installed on your system:

- Information for using the Solaris patch management utility, `patchadd`, is provided in Solaris Administration Guide: Basic Administration at http://download.oracle.com/docs/cd/E19253-01/817-1985/index.html.
- Information for using Solaris Live Upgrade to apply patches is provided in the Solaris installation guide for Live Upgrade and upgrade planning at http://download.oracle.com/docs/cd/E19253-01/821-0438/index.html.

If some patches must be applied when the node is in noncluster mode, you can apply them in a rolling fashion, one node at a time, unless a patch's instructions require that you shut down the entire cluster. Follow procedures in How to Apply a Rebooting Patch (Node) in Oracle Solaris Cluster System Administration Guide to prepare the node and boot it into noncluster mode. For ease of installation, consider applying all patches at once to a node that you place in noncluster mode.

Top

Patch for Cluster Support for Sun StorageTek 2530 Array

The Sun StorageTek Common Array Manager (CAM) software, minimum required Version 6.0.1, provides SCSI3 or PGR support for the Sun StorageTek 2530 array for up to three nodes. The patch is not a required upgrade for the Sun StorEdge 6130, 2540, 6140, and 6540, and StorageTek FLX240, FLX280 and FLX380 platforms. The CAM patch is available from the Sun Download Center.

Top

My Oracle Support

The My Oracle Support Online Web site provides 24-hour access to the most up-to-date information regarding patches, software, and firmware for Oracle products. Access the site at My Oracle Support for the most current matrixes of supported software, firmware, and patch revisions.

Before you install Oracle Solaris Cluster 3.3 software and apply patches to a cluster component (Solaris OS, Oracle Solaris Cluster software, volume manager software, data services software, or disk hardware), review each README file that accompanies the patches that you retrieved. All cluster nodes must have the same patch level for proper cluster operation.

For specific patch procedures and tips on administering patches, see "Patching Oracle Solaris Cluster Software and Firmware" in Oracle Solaris Cluster System Administration Guide.

Top

Product Localization

Localization for certain components of Oracle Solaris Cluster 3.3 software is as follows:

Component	Localization
Software Command Line	Japanese, Simplified Chinese
Software GUI	French, Japanese, Simplified Chinese, Spanish
Online Help	French, Japanese, Simplified Chinese, Spanish
Man Pages	Japanese

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`

Top

Documentation Set

This section lists the collections of the Oracle Solaris Cluster 3.3 user documentation set.

Oracle Solaris Cluster 3.3 Software Manuals

Part Number	Book Title
821-1261	Oracle Solaris Cluster 3.3 Documentation Center
821-1254	Oracle Solaris Cluster Concepts Guide
821-1258	Oracle Solaris Cluster Data Services Developer's Guide
821-1260	Oracle Solaris Cluster Data Services Planning and Administration Guide
821-1259	Oracle Solaris Cluster Error Messages Guide
821-1253	Oracle Solaris Cluster Overview
821-1573	Oracle Solaris Cluster Quick Reference
821-1255	Oracle Solaris Cluster Software Installation Guide
821-1257	Oracle Solaris Cluster System Administration Guide
821-1256	Oracle Solaris Cluster Upgrade Guide

Top

Oracle Solaris Cluster 3.3 Reference Manuals

Part Number	Book Title
821-1263	Oracle Solaris Cluster Reference Manual
821-1264	Oracle Solaris Cluster Data Services Reference Manual

821-1554	Oracle Solaris Cluster Quorum Server Reference Manual

Top

Oracle Solaris Cluster 3.3 Data Service Manuals (SPARC Platform Edition)

Part Number	Book Title
821-1572	Oracle Solaris Cluster Data Service for Agfa IMPAX Guide
821-1571	Oracle Solaris Cluster Data Service for Apache Guide
821-1522	Oracle Solaris Cluster Data Service for Apache Tomcat Guide
821-1523	Oracle Solaris Cluster Data Service for DHCP Guide
821-1524	Oracle Solaris Cluster Data Service for DNS Guide
821-1525	Oracle Solaris Cluster Data Service for Informix Guide
821-1526	Oracle Solaris Cluster Data Service for Kerberos Guide
821-1527	Oracle Solaris Cluster Data Service for MaxDB Guide
821-1528	Oracle Solaris Cluster Data Service for MySQL Guide
821-1932	Oracle Solaris Cluster Data Service for MySQL Cluster Guide
821-1530	Oracle Solaris Cluster Data Service for NFS Guide
821-1531	Oracle Solaris Cluster Data Service for Oracle Guide
821-1532	Oracle Solaris Cluster Data Service for Oracle Application Server Guide
821-2132	Oracle Solaris Cluster Data Service for Oracle Business Intelligence Enterprise Edition Guide
821-1533	Oracle Solaris Cluster Data Service for Oracle E-Business Suite Guide
821-1541	Oracle Solaris Cluster Data Service for Oracle Grid Engine Guide
821-1262	Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide
821-1687	Oracle Solaris Cluster Data Service for Oracle VM Server for SPARC Guide
821-2109	Oracle Solaris Cluster Data Service for PeopleSoft Enterprise Guide
821-1534	Oracle Solaris Cluster Data Service for PostgreSQL Guide
821-1535	Oracle Solaris Cluster Data Service for Samba Guide
821-1536	Oracle Solaris Cluster Data Service for SAP Guide
821-1537	Oracle Solaris Cluster Data Service for SAP liveCache Guide
821-1538	Oracle Solaris Cluster Data Service for SAP Web Application Server Guide
821-1539	Oracle Solaris Cluster Data Service for Siebel Guide
821-1540	Oracle Solaris Cluster Data Service for Solaris Zones Guide
821-1542	Oracle Solaris Cluster Data Service for Sun Java System Application Server EE (HADB) Guide
821-1543	Oracle Solaris Cluster Data Service for Sun Java System Application Server Guide
821-1544	Oracle Solaris Cluster Data Service for Sun Java System Message Queue Guide
821-1546	Oracle Solaris Cluster Data Service for Sun Java System Web Server Guide
821-1547	Oracle Solaris Cluster Data Service for SWIFTAlliance Access Guide
821-1548	Oracle Solaris Cluster Data Service for SWIFTAlliance Gateway Guide

821-1549	Oracle Solaris Cluster Data Service for Sybase ASE Guide
821-1550	Oracle Solaris Cluster Data Service for WebLogic Server Guide
821-1551	Oracle Solaris Cluster Data Service for WebSphere Message Broker Guide
821-1552	Oracle Solaris Cluster Data Service for WebSphere MQ Guide

Top

Oracle Solaris Cluster 3.3 Data Service Manuals (x86 Platform Edition)

Part Number	Book Title
821-1571	Oracle Solaris Cluster Data Service for Apache Guide
821-1522	Oracle Solaris Cluster Data Service for Apache Tomcat Guide
821-1523	Oracle Solaris Cluster Data Service for DHCP Guide
821-1524	Oracle Solaris Cluster Data Service for DNS Guide
821-1525	Oracle Solaris Cluster Data Service for Informix Guide
821-1526	Oracle Solaris Cluster Data Service for Kerberos Guide
821-1527	Oracle Solaris Cluster Data Service for MaxDB Guide
821-1528	Oracle Solaris Cluster Data Service for MySQL Guide
821-1932	Oracle Solaris Cluster Data Service for MySQL Cluster Guide
821-1530	Oracle Solaris Cluster Data Service for NFS Guide
821-1531	Oracle Solaris Cluster Data Service for Oracle Guide
821-1532	Oracle Solaris Cluster Data Service for Oracle Application Server Guide
821-1541	Oracle Solaris Cluster Data Service for Oracle Grid Engine Guide
821-1262	Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide
821-1534	Oracle Solaris Cluster Data Service for PostgreSQL Guide
821-1535	Oracle Solaris Cluster Data Service for Samba Guide
821-1536	Oracle Solaris Cluster Data Service for SAP Guide
821-1537	Oracle Solaris Cluster Data Service for SAP liveCache Guide
821-1538	Oracle Solaris Cluster Data Service for SAP Web Application Server Guide
821-1540	Oracle Solaris Cluster Data Service for Solaris Zones Guide
821-1542	Oracle Solaris Cluster Data Service for Sun Java System Application Server EE (HADB) Guide
821-1543	Oracle Solaris Cluster Data Service for Sun Java System Application Server Guide
821-1544	Oracle Solaris Cluster Data Service for Sun Java System Message Queue Guide
821-1546	Oracle Solaris Cluster Data Service for Sun Java System Web Server Guide
821-1549	Oracle Solaris Cluster Data Service for Sybase ASE Guide
821-1550	Oracle Solaris Cluster Data Service for WebLogic Server Guide
821-1551	Oracle Solaris Cluster Data Service for WebSphere Message Broker Guide
821-1552	Oracle Solaris Cluster Data Service for WebSphere MQ Guide

Top

Oracle Solaris Cluster 3.3 Hardware Collection (SPARC Platform Edition)

Part Number	Book Title
821-1555	Oracle Solaris Cluster 3.3 Hardware Administration Manual
821-1556	Oracle Solaris Cluster 3.3 With Network-Attached Storage Devices Manual
821-1557	Oracle Solaris Cluster 3.3 With SCSI JBOD Storage Device Manual
821-1558	Oracle Solaris Cluster 3.3 With StorageTek RAID Arrays Manual
821-1559	Oracle Solaris Cluster 3.3 With StorageTek 2540 RAID Arrays Manual
821-1560	Oracle Solaris Cluster 3.3 With Sun StorEdge 3310 or 3320 SCSI RAID Array Manual
821-1561	Oracle Solaris Cluster 3.3 With Sun StorEdge 3510 or 3511 FC RAID Array Manual
821-1562	Oracle Solaris Cluster 3.3 With Sun StorEdge 3900 Series or Sun StorEdge 6900 Series System Manual
821-1563	Oracle Solaris Cluster 3.3 With Sun StorEdge 6120 Array Manual
821-1564	Oracle Solaris Cluster 3.3 With Sun StorEdge 6130 Array Manual
821-1565	Oracle Solaris Cluster 3.3 With Sun StorEdge 6320 System Manual
821-1566	Oracle Solaris Cluster 3.3 With Sun StorEdge 9900 Series Storage Device Manual
821-1567	Oracle Solaris Cluster 3.3 With StorEdge A1000 Array, Netra st A1000 Array, or StorEdge A3500 System Manual
821-1568	Oracle Solaris Cluster 3.3 With Fibre Channel JBOD Storage Device Manual
821-1569	Oracle Solaris Cluster 3.3 With Sun StorEdge T3 or T3+ Array Manual
821-1570	Oracle Solaris Cluster 3.3 With Sun StorEdge 3500FC System Manual

Top

Oracle Solaris Cluster 3.3 Hardware Collection (x86 Platform Edition)

Part Number	Book Title
821-1555	Oracle Solaris Cluster 3.3 Hardware Administration Manual
821-1556	Oracle Solaris Cluster 3.3 With Network-Attached Storage Devices Manual
821-1557	Oracle Solaris Cluster 3.3 With SCSI JBOD Storage Device Manual
821-1558	Oracle Solaris Cluster 3.3 With StorageTek RAID Arrays Manual
821-1559	Oracle Solaris Cluster 3.3 With StorageTek 2540 RAID Arrays Manual
821-1561	Oracle Solaris Cluster 3.3 With Sun StorEdge 3510 or 3511 FC RAID Array Manual
821-1563	Oracle Solaris Cluster 3.3 With Sun StorEdge 6120 Array Manual
821-1564	Oracle Solaris Cluster 3.3 With Sun StorEdge 6130 Array Manual
821-1565	Oracle Solaris Cluster 3.3 With Sun StorEdge 6320 System Manual
821-1566	Oracle Solaris Cluster 3.3 With Sun StorEdge 9900 Series Storage Device Manual

Top

Documentation Issues

This section discusses errors or omissions for documentation, online help, or man pages in the Oracle Solaris Cluster 3.3 release.

- System Administration Guide
- Data Service for Apache Tomcat Guide
- Data Service for Oracle Guide
- Data Service for Oracle Real Application Clusters Guide
- Network-Attached Storage Manual
- Man Pages
- Online Help

Top

Data Service for Oracle Real Application Clusters Guide

This section discusses errors and omissions in the Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide.

CLI Support to Configure Oracle RAC Instance Proxy Resource Group

Support is added to the Oracle Solaris Cluster maintenance commands to configure an Oracle Real Application Clusters (Oracle RAC) instance proxy resource group when Oracle ASM is not also configured. This support is delivered starting in the following HA-Oracle patch versions:

- SPARC: 145335-01
- x86: 145336-01

For supporting configuration instructions, see Configuring an Oracle RAC Instance Proxy Resource Group Without ASM (CLI) .

Procedures for removing the Grid Infrastructure resource are also provided. See Removing an Oracle Grid Infrastructure Resource.

Top

Configuring an Oracle RAC Instance Proxy Resource Group Without Oracle ASM (CLI)

For Oracle 11g release 2, a patch and additional configuration procedures are required to support using Oracle Solaris Cluster maintenance commands to configure an Oracle RAC instance proxy resource group without Oracle ASM. Perform the following tasks to add this support:

Task	Additional Information
1. Apply at least the following minimum HA-Oracle patch version on each node that will run Oracle 11g release 2: • SPARC: 145335-01 • x86: 145336-01	
2. Configure the disk set or disk group to use.	Follow one of the following procedures in the Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide: • If you are using Solaris Volume Manager for Sun Cluster, create a multi-owner disk set. Follow procedures in How to Create a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster for the Oracle RAC Database . • If you are using Veritas Volume Manager with the cluster feature, create a shared disk group. Follow procedures in How to Create a VxVM Shared-Disk Group for the Oracle RAC Database .

3. Perform procedures in Creating an Oracle Grid Infrastructure Resource.	When you configure Oracle RAC 11g release 2, this procedure replaces How to Create an Oracle Clusterware Resource for Interoperation With Sun Cluster in Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide. You can alternatively perform this task by using the Oracle Solaris Cluster Manager GUI or `clsetup` wizards, as described in How to Enable Oracle Solaris Cluster and Oracle 10g Release 2 or 11g Oracle Clusterware to Interoperate . If you do, and a patch for CR 6995923 is not yet available, you must first perform Step 1 through Step 3 of this procedure, to create the prerequisite Oracle Grid Infrastructure resource, before you run the wizard. The wizard does not currently create a missing Oracle Grid Infrastructure resource for you.
4. Perform the appropriate CLI procedure in Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide to create Oracle Solaris Cluster resources for interoperation with Oracle .	When an offline-restart dependency is set, the VALIDATE method of the `SUNW.scalable_rac_server_proxy` resource transparently modifies the appropriate Grid Infrastructure database resource as follows: • `ora.dbname.db` --> [add hard-start dependency] --> `sun.scal-dg1-rs` • `ora.dbname.db` --> [add hard-start dependency] --> `sun.scal-dg2-rs`

Top

Creating an Oracle Grid Infrastructure Resource

Perform this procedure to manually create an Oracle Grid Infrastructure resource that proxies the Oracle Solaris Cluster `SUNW.ScalDeviceGroup` or `SUNW.ScalMountPoint` resource. In this example procedure, the Oracle Grid Infrastructure resource is named `sun.`*`resource`*. You configure `sun.`*`resource`* to ensure that the corresponding Oracle database is not started until `sun.`*`resource`* is online. The `sun.`*`resource`* resource comes online only if the corresponding `SUNW.ScalDeviceGroup` or `SUNW.ScalMountPoint` resource is online. And the `SUNW.ScalDeviceGroup` or `SUNW.ScalMountPoint` resource only comes online if the actual volume-manager disk set or disk group or the mount point is online.

Top

How to Create an Oracle Grid Infrastructure Resource

Perform this procedure on one node of the cluster.

1. Become superuser.

2. Create the Oracle Grid Infrastructure **`sun.storage_proxy.type`** resource type.

In this example procedure, the `/export/home/oinstall/app/11.2.0/grid` directory is the Oracle Grid Infrastructure home directory.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl \
add type sun.storage_proxy.type \
-basetype cluster_resource \
-attr \
"ATTRIBUTE=ACTION_SCRIPT,TYPE=string", \
"ATTRIBUTE=HOSTING_MEMBERS,TYPE=string", \
"ATTRIBUTE=CARDINALITY,TYPE=string", \
"ATTRIBUTE=PLACEMENT,TYPE=string", \
"ATTRIBUTE=SCRIPT_TIMEOUT,TYPE=int", \
"ATTRIBUTE=RESTART_ATTEMPTS,TYPE=int", \
"ATTRIBUTE=ACL,TYPE=string", \
"ATTRIBUTE=VERSION,TYPE=string"
```

3. Create the Oracle Solaris Cluster **`sun.`*`resource`*** resource of type **`sun.storage_proxy.type`**.

The Oracle Grid Infrastructure resource name uses the form `sun.`*`sc-resource`*, where *`sc-resource`* is the name of the `SUNW.ScalDeviceGroup` or `SUNW.ScalMountPoint` resource. For example, the Oracle Solaris Cluster resource `scal-dg1-rs` is identified as `sun.scal-dg1-rs` in Oracle Grid Infrastructure.

You must also set the following attributes to values that are appropriate for your cluster:

- `HOSTING_MEMBERS` = The list of nodes in the cluster membership
- `CARDINALITY` = The number of nodes in the cluster membership

To create the Oracle Grid Infrastructure `storage_proxy` resource, execute the following command, substituting appropriate values for your cluster:

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl add resource sun.scal-dg1-rs \
-type sun.storage_proxy.type \
-attr "ACTION_SCRIPT='/opt/SUNWscor/dsconfig/bin/scproxy_crs_action' \
CARDINALITY='3' \
SCRIPT_TIMEOUT='20' \
PLACEMENT='restricted' \
RESTART_ATTEMPTS='60' \
HOSTING_MEMBERS='pnsx1 pnsx2 pnsx3' \
VERSION='1' "
```

4. Determine the DBA group of the Oracle Grid Infrastructure installation.

In this example procedure, the `/export/home/oinstall/app/11.2.0/grid` directory is the Oracle Grid Infrastructure home directory.

```
# echo `/export/home/oinstall/app/11.2.0/grid/bin/osdbagrp`
griddba
```

5. Set the primary group of the Oracle Grid Infrastructure **`storage_proxy`** resource to the group determined in Step 4.

In this example procedure, the `/export/home/oinstall/app/11.2.0/grid` directory is the Oracle Grid Infrastructure home directory.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl setperm resource sun.scal-dg1-rs -g "griddba"
```

6. Determine the DBA group of the Oracle Database Software installation.

In this example procedure, the `/export/home/oinstall/app/oracle/product/11.2.0/db` directory is the Oracle Database Software home directory.

```
# echo `/export/home/oinstall/app/oracle/product/11.2.0/db/bin/osdbagrp`
dba
```

7. Set the group permissions of the Oracle Grid Infrastructure **`storage_proxy`** resource to the group determined in Step 6.

Omit this step if the Oracle Grid Infrastructure installation DBA group determined in Step 4 and the Oracle Database Software installation DBA group determined in Step 6 are the same DBA group.

In this example procedure, the `/export/home/oinstall/app/11.2.0/grid` directory is the Oracle Grid Infrastructure home directory.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl setperm resource sun.scal-dg1-rs -u "group:dba:r-x"
```

8. Bring online the Oracle Grid Infrastructure **`storage_proxy`** resource.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl start resource sun.scal-dg1-rs
CRS-2672: Attempting to start 'sun.scal-dg1-rs' on 'pnsx2'
CRS-2672: Attempting to start 'sun.scal-dg1-rs' on 'pnsx3'
CRS-2672: Attempting to start 'sun.scal-dg1-rs' on 'pnsx1'
CRS-2676: Start of 'sun.scal-dg1-rs' on 'pnsx2' succeeded
CRS-2676: Start of 'sun.scal-dg1-rs' on 'pnsx3' succeeded
CRS-2676: Start of 'sun.scal-dg1-rs' on 'pnsx1' succeeded
```

Top

Removing an Oracle Grid Infrastructure Resource

This section contains the following procedures to remove an Oracle Grid Infrastructure resource.

- How to Remove a Dependency
- How to Delete the `sun.resource` Resource

Top

How to Remove a Dependency

This procedure shows how to set the offline-restart dependency to remove a dependency.

1. Become superuser.

2. Display the current start dependency that the database has on the Oracle Grid Infrastructure **`storage_proxy`** resource.

In this example procedure, the `/export/home/oinstall/app/11.2.0/grid` directory is the Oracle Grid Infrastructure home directory, the database name is `testdb`, and the Oracle Grid Infrastructure `storage_proxy` resource name is `sun.scal-dg1-rs`.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl stat res ora.testdb.db -p | grep START_DEPENDENCIES
START_DEPENDENCIES=hard(sun.scal-dg1-rs)
weak(type:ora.listener.type,global:type:ora.scan_listener.type,uniform:ora.ons,uniform:ora.eons)

# clresource show -p resource_dependencies_offline_restart rac-server-proxy-rs
=== Resources ===

Resource: rac-server-proxy-rs
Resource_dependencies_offline_restart: crs-fw-rs scal-dg1-rs

— Standard and extension properties —
```

3. Remove the offline-restart dependency on the **`SUNW.ScalDeviceGroup`** or **`SUNW.ScalMountPoint`** resource from the Oracle RAC instance proxy resource.

This command clears the dependencies that the Oracle Grid Infrastructure database resource has on the Oracle Grid Infrastructure `storage_proxy` resource. Note that the command includes the minus (-) symbol.

```
# clresource set -p resource_dependencies_offline_restart-=scal-dg1-rs rac-server-proxy-rs
```

4. Verify that the start dependency on the Oracle Grid Infrastructure resource is removed.

In this example procedure, the `/export/home/oinstall/app/11.2.0/grid` directory is the Oracle Grid Infrastructure home directory.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl stat res ora.testdb.db -p | grep START_DEPENDENCIES
START_DEPENDENCIES=weak(type:ora.listener.type,global:type:ora.scan_listener.type,uniform:ora.ons,uniform:ora.eons

# clresource show -p resource_dependencies_offline_restart rac-server-proxy-rs

=== Resources ===

Resource: rac-server-proxy-rs
Resource_dependencies_offline_restart: crs-fw-rs

— Standard and extension properties —
```

Top

How to Delete the *`sun.resource`* Resource

1. Become superuser.

2. Ensure that the dependency is removed, as described in How to Remove a Dependency, and that **sun.resource** is stopped.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl stop res sun.scal-dg1-rs
CRS-2673: Attempting to stop 'sun.scal-dg1-rs' on 'pnsx3'
CRS-2673: Attempting to stop 'sun.scal-dg1-rs' on 'pnsx1'
CRS-2673: Attempting to stop 'sun.scal-dg1-rss' on 'pnsx2'
CRS-2677: Stop of 'sun.scal-dg1-rs' on 'pnsx3' succeeded
CRS-2677: Stop of 'sun.scal-dg1-rs' on 'pnsx1' succeeded
CRS-2677: Stop of 'sun.scal-dg1-rs' on 'pnsx2' succeeded
```

3. Delete **sun.*resource***.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl delete res sun.scal-dg1-rs
```

4. Verify that **sun.*resource*** is deleted.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl stat res sun.scal-dg1-rs -p
CRS-210: Could not find resource 'sun.scal-dg1-rs'.
```

Top

System Administration Guide

Extraneous Information in "Adding a Quorum Device"

The section "Adding a Quorum Device" refers to the use of Oracle's Sun Storage 7000 Unified Storage System as a NAS quorum device. Support for Oracle's Sun Storage 7000 Unified Storage System as a NAS device has not been enabled in the 3.3 release. Ignore the information.

Top

Data Service for Apache Tomcat Guide

Changes With Apache Tomcat 5.5.28 and 6.0.29

The following information additions and changes support the `wget` probe algorithm for Oracle Solaris Cluster HA for Apache Tomcat (HA for Apache Tomcat) using Apache Tomcat versions starting with 5.5.28 and 6.0.29. This functionality is delivered in the following Oracle Solaris Cluster 3.3 patches for HA for Apache Tomcat:

- SPARC: 145341-01
- x86: 145342-01

- New Restriction For Apache Tomcat 5.5.28 and 6.0.29
- New Requirement for Apache Tomcat 5.5.28 and 6.0.29
- Installing and Configuring Apache Tomcat
- How to Register and Configure Solaris Cluster HA for Apache Tomcat as a Failover Data Service
- How to Register and Configure Solaris Cluster HA for Apache Tomcat as a Multiple Masters Data Service
- How to Register and Configure Solaris Cluster HA for Apache Tomcat as a Scalable Data Service
- Structure of the Apache Tomcat HA parameter file
- Strategy to Chose the TestCmd and the ReturnString Variable
- Probing Algorithm and Functionality
- Deployment Example: Installing Apache Tomcat in the Global Zone
- Example: Modifying the Apache Tomcat Configuration
- Deployment Example: Installing Apache Tomcat in a Failover Zone
- Example: Installing the Apache Tomcat Software on Shared Storage
- Deployment Example: Installing Apache Tomcat in a Non-Global Zone
- Example: Modifying the Apache Tomcat Configuration File

New Restriction For Apache Tomcat 5.5.28 and 6.0.29

Starting with Apache Tomcat versions 5.5.28 and 6.0.29, you must use the `wget` probe algorithm. This bypasses the normal probe, enabling the TestUrl parameter to work correctly. You can optionally use the `wget` probe algorithm for earlier versions as well.

Top

New Requirement for Apache Tomcat 5.5.28 and 6.0.29

To support the `wget` probe algorithm, the `wget` command must be available from every node of the cluster.

Top

Installing and Configuring Apache Tomcat

In the bullet list of considerations to determine how Apache Tomcat will be deployed with Oracle Solaris Cluster, you must also determine which Apache Tomcat version to deploy.

Top

How to Register and Configure Solaris Cluster HA for Apache Tomcat as a Failover Data Service

The following replaces Step 1:

1. Prepare the parameter file, which is required by HA for Apache Tomcat.

```
# cd /opt/SUNWsctomcat/bin
# cp pfile desired-place
```

Edit the parameter file `pfile` and follow the comments within that file. For example:

```
#!/usr/bin/ksh
#
# CDDL HEADER START
#
# The contents of this file are subject to the terms of the
# Common Development and Distribution License (the License).
# You may not use this file except in compliance with the License.
#
# You can obtain a copy of the license at usr/src/CDDL.txt
# or http://www.opensolaris.org/os/licensing.
# See the License for the specific language governing permissions
# and limitations under the License.
#
# When distributing Covered Code, include this CDDL HEADER in each
# file and include the License file at usr/src/CDDL.txt.
# If applicable, add the following below this CDDL HEADER, with the
# fields enclosed by brackets \[\] replaced with your own identifying
# information: Portions Copyright \[yyyy\] \[name of copyright owner\]
#
# CDDL HEADER END
#
#
# Copyright (c) 2006, 2010, Oracle and/or its affiliates. All rights reserved.
#

#ident  "@(#)pfile.ksh  1.6     10/08/16"

# Set the Apache Tomcat specific environment variables which the start, stop
# and check fuctions will use
#
# EnvScript     Script to set runtime environment for tomcat
# User          Apache Tomcat User
# Basepath      Absolute path to Tomcat Home directory i.e. 4.x CATALINA_HOME
#               or TOMCAT_HOME for 3.x
# Host          Hostname to test Apache Tomcat
# Port          Port where Apache Tomcat is configured
# TestCmd       Apache Tomcat test command
#
# Some versions of Apache Tomcat do not react on the probe algorithm with the
# http get command. Because of this we provide an alternative probe method
# which uses wget. If wget is not installed on your system you must install it
# manually and provide the WgetPath variable. If the WgetPath variable is not set,
# it defaults to /usr/sfw/bin/wget.
# You can use the wget algorithm for https probe as well which are not possible
# with the http get command.
#
# Use the variables TestUrl and WgetPath as a replacement for:
# Port, Host, TestCmd
# You can specify only one set of variables.
#
# TestUrl       Url where Tomcat serve a web site. This can be done either
#               by the http or https  protocol
#               Format: http://hostname:port/startpage
#               Example: http://localhost:8080/
#               You might want to include options here like:
#               "--no-cookies http://localhost:8080/"
#               The requirement here is that wget -O - ${TestUrl} produces
#               the output containing the ReturnString.
# WgetPath      Abolute pat to wget, leaving this variable empty defaults to:
#               /usr/sfw/bin/wget
# Startwait     Sleeping $Startwait seconds after completion of the
#               start command
# ReturnString  This string must be present in the output of the http get command
#               or in the output of the wget command.
EnvScript=
User=
Basepath=
Host=
Port=8080
TestCmd="get /index.jsp"
TestUrl=
WgetPath=
ReturnString="CATALINA"
Startwait=20
```

The following is an example for Apache Tomcat 4.1.24:

```
EnvScript=/global/mnt1/tomtest/env.ksh
User=tomtest
Basepath=/global/mnt1/tomtest/jakarta-tomcat-4.1.24
Host=tomcat-lh
Port=8080
TestCmd="get /index.jsp"
TestUrl=
WgetPath=
ReturnString="CATALINA"
Startwait=20
```

This example for Apache Tomcat 4.1.24 configures a standalone web server with the default start page `index.jsp`. The start page contains the string `CATALINA`.

To use the alternative `wget` probe, assuming `wget` is available under `/usr/sfw/bin/wget`, instead set the following variables:

```
EnvScript=/global/mnt1/tomtest/env.ksh
User=tomtest
Basepath=/global/mnt1/tomtest/jakarta-tomcat-4.1.24
Host=
Port=
TestCmd=
TestUrl="http://tomcat-lh:8080/"
WgetPath=
ReturnString="CATALINA"
Startwait=20
```

Note
The parameter files must be available on every node that can host the Apache Tomcat data service. For a scalable or multiple-masters configuration, store the parameter files on the local file system of every node or non-global zone or on the shared storage. The parameter files must not differ for any instance of Apache Tomcat on the various nodes.

Repeat Step 1 for every Apache Tomcat instance you need.

Top

How to Register and Configure Solaris Cluster HA for Apache Tomcat as a Multiple Masters Data Service

The following replaces Step 2:

2. Prepare the parameter file, which is required by HA for Apache Tomcat.

```
# cd /opt/SUNWsctomcat/util
# cp sctomcat_config desired-place
```

Edit the parameter file `pfile` and follow the comments within that file. For example:

```
#!/usr/bin/ksh
#
# CDDL HEADER START
#
# The contents of this file are subject to the terms of the
# Common Development and Distribution License (the License).
# You may not use this file except in compliance with the License.
#
# You can obtain a copy of the license at usr/src/CDDL.txt
# or http://www.opensolaris.org/os/licensing.
# See the License for the specific language governing permissions
# and limitations under the License.
#
```

```
# When distributing Covered Code, include this CDDL HEADER in each
# file and include the License file at usr/src/CDDL.txt.
# If applicable, add the following below this CDDL HEADER, with the
# fields enclosed by brackets \[\] replaced with your own identifying
# information: Portions Copyright \[yyyy\] \[name of copyright owner\]
#
# CDDL HEADER END
#

#
# Copyright (c) 2006, 2010, Oracle and/or its affiliates. All rights reserved.
#

#ident  "@(#)pfile.ksh  1.6     10/08/16"

# Set the Apache Tomcat specific environment variables which the start, stop
# and check fuctions will use
#
# EnvScript     Script to set runtime environment for tomcat
# User          Apache Tomcat User
# Basepath      Absolute path to Tomcat Home directory i.e. 4.x CATALINA_HOME
#               or TOMCAT_HOME for 3.x
# Host          Hostname to test Apache Tomcat
# Port          Port where Apache Tomcat is configured
# TestCmd       Apache Tomcat test command
#
# Some versions of Apache Tomcat do not react on the probe algorithm with the
# http get command. Because of this we provide an alternative probe method
# which uses wget. If wget is not installed on your system you must install it
# manually and provide the WgetPath variable. If the WgetPath variable is not set,
# it defaults to /usr/sfw/bin/wget.
# You can use the wget algorithm for https probe as well which are not possible
# with the http get command.
#
# Use the variables TestUrl and WgetPath as a replacement for:
# Port, Host, TestCmd
# You can specify only one set of variables.
#
# TestUrl       Url where Tomcat serve a web site. This can be done either
#               by the http or https  protocol
#               Format: http://hostname:port/startpage
#               Example: http://localhost:8080/
#               You might want to include options here like:
#               "--no-cookies http://localhost:8080/"
#               The requirement here is that wget -O - ${TestUrl} produces
#               the output containing the ReturnString.
# WgetPath      Abolute pat to wget, leaving this variable empty defaults to:
#               /usr/sfw/bin/wget
# Startwait     Sleeping $Startwait seconds after completion of the
#               start command
# ReturnString  This string must be present in the output of the http get command
#               or in the output of the wget command.
EnvScript=
User=
Basepath=
Host=
Port=8080
TestCmd="get /index.jsp"
TestUrl=
WgetPath=
```

```
ReturnString="CATALINA"
Startwait=20
```

The following is an example for Apache Tomcat 4.1.24.

```
EnvScript=/tomcat/env.ksh
User=tomcat
Basepath=/tomcat/jakarta-tomca-4.1.24
Host=localhost
Port=8080
TestCmd="get /index.jsp"
TestUrl=
WgetPath=
ReturnString="CATALINA"
Startwait=20
```

This example for Apache Tomcat 4.1.24 configures a standalone web server with the default start page `index.jsp`. The start page contains the string `CATALINA`.

Depending on the selected probe method, the Host parameter or the IP alias in the TestUrl parameter depends on the location of the `pfile`:

- If the `pfile` is stored on the shared storage, you must use `localhost`.
- If the `pfile` is stored on the local storage, it can be either `localhost` or the node's hostname.

To use the alternative `wget` probe, assuming `wget` is available under `/usr/sfw/bin/wget`, set the following variables:

```
EnvScript=/global/mnt1/tomtest/env.ksh
User=tomtest
Basepath=/global/mnt1/tomtest/jakarta-tomcat-4.1.24
Host=
Port=
TestCmd=
TestUrl=”http://localhost:8080/”
WgetPath=
ReturnString="CATALINA"
Startwait=20
```

Note
The parameter files must be available on every node that can host the Apache Tomcat data service. For a scalable or multiple-masters configuration, store the parameter files on the local file system of every node or non-global zone or on the shared storage. The parameter files must not differ for any instance of Apache Tomcat on the various nodes.

Repeat this Step for every Apache Tomcat instance and every node you need.

Top

How to Register and Configure Solaris Cluster HA for Apache Tomcat as a Scalable Data Service

The following replaces Step 2:

2. Prepare the parameter file, which is required by Oracle Solaris Cluster HA for Apache Tomcat.

```
# cd /opt/SUNWsctomcat/bin
# cp pfile desired-place
```

Edit the parameter file `pfile` and follow the comments within that file. For example:

```
#!/usr/bin/ksh
#
# CDDL HEADER START
```

```
#
# The contents of this file are subject to the terms of the
# Common Development and Distribution License (the License).
# You may not use this file except in compliance with the License.
#
# You can obtain a copy of the license at usr/src/CDDL.txt
# or http://www.opensolaris.org/os/licensing.
# See the License for the specific language governing permissions
# and limitations under the License.
#
# When distributing Covered Code, include this CDDL HEADER in each
# file and include the License file at usr/src/CDDL.txt.
# If applicable, add the following below this CDDL HEADER, with the
# fields enclosed by brackets \[\] replaced with your own identifying
# information: Portions Copyright \[yyyy\] \[name of copyright owner\]
#
# CDDL HEADER END
#

#
# Copyright (c) 2006, 2010, Oracle and/or its affiliates. All rights reserved.
#

#ident  "@(#)pfile.ksh  1.6     10/08/16"

# Set the Apache Tomcat specific environment variables which the start, stop
# and check fuctions will use
#
# EnvScript     Script to set runtime environment for tomcat
# User          Apache Tomcat User
# Basepath      Absolute path to Tomcat Home directory i.e. 4.x CATALINA_HOME
#               or TOMCAT_HOME for 3.x
# Host          Hostname to test Apache Tomcat
# Port          Port where Apache Tomcat is configured
# TestCmd       Apache Tomcat test command
#
# Some versions of Apache Tomcat do not react on the probe algorithm with the
# http get command. Because of this we provide an alternative probe method
# which uses wget. If wget is not installed on your system you must install it
# manually and provide the WgetPath variable. If the WgetPath variable is not set,
# it defaults to /usr/sfw/bin/wget.
# You can use the wget algorithm for https probe as well which are not possible
# with the http get command.
#
# Use the variables TestUrl and WgetPath as a replacement for:
# Port, Host, TestCmd
# You can specify only one set of variables.
#
# TestUrl       Url where Tomcat serve a web site. This can be done either
#               by the http or https  protocol
#               Format: http://hostname:port/startpage
#               Example: http://localhost:8080/
#               You might want to include options here like:
#               "--no-cookies http://localhost:8080/"
#               The requirement here is that wget -O - ${TestUrl} produces
#               the output containing the ReturnString.
# WgetPath      Abolute pat to wget, leaving this variable empty defaults to:
#               /usr/sfw/bin/wget
# Startwait     Sleeping $Startwait seconds after completion of the
#               start command
# ReturnString  This string must be present in the output of the http get command
#               or in the output of the wget command.
EnvScript=
User=
Basepath=
Host=
Port=8080
TestCmd="get /index.jsp"
TestUrl=
WgetPath=
```

```
ReturnString="CATALINA"
Startwait=20
```

The following is an example for Apache Tomcat 4.1.24.

```
EnvScript=/tomcat/env.ksh
User=tomcat
Basepath=/tomcat/jakarta-tomca-4.1.24
Host=localhost
Port=8080
TestCmd="get /index.jsp"
TestUrl=
WgetPath=
ReturnString="CATALINA"
Startwait=20
```

To use the alternative `wget` probe, assuming `wget` is available under `/usr/sfw/bin/wget`, set the following variables:

```
EnvScript=/global/mnt1/tomtest/env.ksh
User=tomtest
Basepath=/global/mnt1/tomtest/jakarta-tomcat-4.1.24
Host=
Port=
TestCmd=
TestUrl="http://localhost:8080/"
WgetPath=
ReturnString="CATALINA"
Startwait=20
```

This example for Apache Tomcat 4.1.24 configures a standalone web server with the default start page `index.jsp`. The start page contains the string `CATALINA`.

Depending on the selected probe method, the Host parameter or the IP alias in the TestUrl parameter depends on the location of the `pfile`.

- If the `pfile` is stored on the shared storage, you must use `localhost`.
- If the `pfile` is stored on the local storage, it can be either `localhost` or the node's hostname.

Note
Do not use the SharedAddress here.

Note
The parameter files must be available on every node that can host the Apache Tomcat data service. For a scalable or multiple-masters configuration, store the parameter files on the local file system of every node or on the shared storage. The parameter files must not differ for an instance of Apache Tomcat on the various nodes.

Repeat this Step for every Apache Tomcat instance and every node you need.

Top

Structure of the Apache Tomcat HA Parameter File

The following are additions to Table 6: Structure of the Apache Tomcat HA parameter file.

Variable	Explanation
TestUrl	This variable presents a complete `http` or `https` URL, which is passed to the `wget` command to test the Tomcat server's sanity.
WgetPath	This variable presents the absolute path to the `wget` binary, if the `wget` binary is not located under `/usr/sfw/bin/wget`.

In addition, TestUrl and WgetPath are added to the list of parameters that are used for starting and stopping Apache Tomcat.

Top

Strategy to Choose the TestCmd and the ReturnString Variable

The following replaces the section "Strategy to Chose the TestCmd and the ReturnString Variable".

Strategy to Choose the TestCmd or TestUrl Variable and the ReturnString Variable

The following alternatives are available to choose the ReturnString variable:

- Take the start page of your application and set the TestCmd to get */start_page*. If you picked the `wget` probe method, set TestUrl to `http://`*`ipalias`*`:`*`port`*`/start_page`. With this strategy, you are monitoring that Apache Tomcat is serving your application.
- Take the Apache Tomcat default start page and set the TestCmd to get */default-startpage*. If you picked the `wget` method, set TestUrl to `http://`*`ipalias`*`:`*`port`*`/`*`default-startpage`*. In this case, set the ReturnString to a string contained in the startpage. This string depends on the deployed Apache Tomcat version; for 3.x it is Tomcat home and for 4.x and 5.x it is CATALINA. With this strategy, you are monitoring that Apache Tomcat is serving its default application.
- Deploy a test application (which is not provided with the HA for Apache Tomcat agent) to Apache Tomcat. Set the TestCmd to get *`startpage_of_the_application`*. If you use the `wget` method, set TestUrl to `http://`*`ipalias`*`:`*`port`*`/`*`startpage_of_the_application`*. In this case, set the ReturnString to a string contained in the start page. With this strategy, you are monitoring that Apache Tomcat is serving your test application.
- If none of above is appropriate, set the TestCmd to get */a-page-that-does-not-exist*. In this case, set the ReturnString to a string contained in the Error Page. With this strategy, you are monitoring that Apache Tomcat is operational, because it registers that it must deliver a page that does not exist.

You can evaluate the different pages by connecting from a browser to *`hostname`*`:`*`port`* and specifying the different pages.

Top

Probing Algorithm and Functionality

In the steps that describe how the sanity of Apache Tomcat is monitored, the following replaces Step 2:

2. Pings the Host, which is configured in the Apache Tomcat HA parameter file, unless the TestUrl parameter is specified in the Tomcat parameter file.

Top

Deployment Example: Installing Apache Tomcat in the Global Zone

In the section "Software Configuration", the following is an additional expectation of the example configuration:

- The `wget` binary is available under `/usr/sfw/bin/wget`.

Top

Example: Modifying the Apache Tomcat Configuration

The following is an addition to Step 4:

As an alternative, you can use the following set of variables:

```
EnvScript=/global/mnt3/tomcat/env.ksh
User=tomcat
Basepath=/global/mnt3/tomcat/apache-tomcat-5.5.17
TestUrl="http://ha-host-1:8080/"
ReturnString="CATALINA"
Startwait=20
```

Top

Deployment Example: Installing Apache Tomcat in a Failover Zone

In the section "Software Configuration", the following is an additional expectation of the example configuration:

- The `wget` binary is available under `/usr/sfw/bin/wget`.

Top

Example: Installing the Apache Tomcat Software on Shared Storage

The following is an addition to Step 6:

As an alternative, you can use the following set of variables:

```
EnvScript=/global/mnt3/tomcat/env.ksh
User=tomcat
Basepath=/global/mnt3/tomcat/apache-tomcat-5.5.17
TestUrl="http://ha-host-1:8080/"
ReturnString="CATALINA"
Startwait=20
```

Top

Deployment Example: Installing Apache Tomcat in a Non-Global Zone

In the section "Software Configuration", the following is an additional expectation of the example configuration:

- The `wget` binary is available under `/usr/sfw/bin/wget`.

Top

Example: Modifying the Apache Tomcat Configuration File

The following is an addition to Step 3:

As an alternative, you can use the following set of variables:

```
EnvScript=/global/mnt3/tomcat/env.ksh
User=tomcat
Basepath=/global/mnt3/tomcat/apache-tomcat-5.5.17
TestUrl="http://ha-host-1:8080/"
ReturnString="CATALINA"
Startwait=20
```

Top

Data Service for Oracle Guide

This section discusses errors and omissions in the Oracle Solaris Cluster Data Service for Oracle Guide.

Support for Oracle Grid Infrastructure for Clusters SCAN Listener with 11g Release 2

When the HA for Oracle data service is deployed using Oracle 11g release 2 Grid Infrastructure for Clusters, the preferred listener for a single-instance database is the Grid Infrastructure for Clusters Single Client Access Name (SCAN) Listener. To leverage SCAN in the HA-Oracle configuration, you must modify the Oracle 11g release 2 single-instance database remote_listener parameter to reflect the SCAN name and port number.

The following example shows the commands and sample output to configure the SCAN name and port.

```
# srvctl config scan
SCAN name: freak-scan-lh, Network: 1/10.11.188.0/255.255.255.0/bge0
SCAN VIP name: scan1, IP: /vzfreak1a/10.11.188.60

# srvctl config scan_listener
SCAN Listener LISTENER_SCAN1 exists. Port: TCP:1521
#

SQL> show parameters listener

NAME          TYPE   VALUE
----------------------------------- ----------- ------------------------------
listener_networks       string
local_listener        string
(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=freak-1)(PORT=1521))))
remote_listener         string


SQL>
SQL> alter system set remote_listener = 'freak-scan-lh:1521' scope=both;

System altered.
SQL> show parameters listener

NAME          TYPE   VALUE
----------------------------------- ----------- ------------------------------
listener_networks       string
local_listener        string
(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=freak-1)(PORT=1521))))
remote_listener         string      freak-scan-lh:1521
SQL>
```

If you do not want to use the SCAN listener, use the `SUNW.oracle_listener` resource type instead and configure it as usual.

Top

Support for Third-Party Volume Managers with Clustered Oracle ASM Disk Groups

For Oracle 11g release 2, a patch and additional configuration procedures are required to support third-party volume managers with clustered Oracle ASM disk groups. Perform the following tasks to add this support:

1. Apply the following minimum HA-Oracle patch on each node that will run Oracle 11g release 2 with clustered Oracle ASM:
   - SPARC: 145335-01
   - x86: 145336-01
2. Configure the disk set or disk group for Oracle ASM to use.
   Follow the appropriate procedure in the Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide:
   - If you are using Solaris Volume Manager for Sun Cluster, create a multi-owner disk set. Follow procedures in "How to Create a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster for the Oracle RAC Database" .
   - If you are using Veritas Volume Manager with the cluster feature, create a shared disk group. Follow procedures in "How to Create a VxVM Shared-Disk Group for the Oracle RAC Database" .
3. Perform procedures in "Creating an Oracle Grid Infrastructure Resource" .

Procedures for removing the Grid Infrastructure resource are also provided. See "Removing an Oracle Grid Infrastructure Resource" .

Top

Creating an Oracle Grid Infrastructure Resource

Perform this procedure to manually create a Grid Infrastructure resource that proxies the Oracle Solaris Cluster `SUNW.ScalDeviceGroup` resource. In this example procedure, the Grid Infrastructure resource is named `sun.`*`resource`*. You configure `sun.`*`resource`* to ensure that the corresponding Oracle ASM disk group is not mounted until `sun.`*`resource`* is online. The `sun.`*`resource`* resource comes online only if the corresponding `SUNW.ScalDeviceGroup` resource is online. And the `SUNW.ScalDeviceGroup` resource only comes online if the actual volume-manager disk set or disk group is online.

To ensure that the Oracle ASM disk group benefits from this dependency chain, after you define `sun.`*`resource`*, you modify the

appropriate Grid Infrastructure Oracle ASM disk-group resource so that the hard-start dependency includes `sun.resource`. Modifying the hard-start dependency of the Grid Infrastructure Oracle ASM disk-group resource can only be performed by the `SUNW.scalable_asm_diskgroup_proxy` resource by using the VALIDATE method. Therefore, you must set an offline-restart dependency between the `SUNW.scalable_asm_diskgroup_proxy` and `SUNW.ScalDeviceGroup` resources.

Top

How to Create an Oracle Grid Infrastructure Resource

This example procedure sets the following dependencies:

- A Grid Infrastructure resource, `sun.resource`, where `resource` corresponds to an Oracle Solaris Cluster `SUNW.ScalDeviceGroup` resource name:
  - `sun.scal-asmdg1-rs`
  - `sun.scal-asmdg2-rs`
- Oracle Solaris Cluster offline-restart dependencies:
  - `asm-data1-rs` --> [add offline-restart dependency] -> `scal-asmdg1-rs`
  - `asm-data2-rs` --> [add offline-restart dependency] -> `scal-asmdg2-rs`

When an offline-restart dependency is set, the VALIDATE method of the SUNW.scalable_asm_diskgroup_proxy resource will transparently modify the appropriate Grid Infrastructure ASM disk group resource as follows:

- Grid Infrastructure hard-start dependencies:
  - `ora.data1.dg` -> [add hard-start dependency] -> `sun.scal-asmdg1-rs`
  - `ora.data1.dg` -> [add hard-start dependency] -> `sun.scal-asmdg1-rs`

Perform this procedure on one node of the cluster.

1. Become superuser.
2. Create the **`sun.storage_proxy.type`** resource type.
   In this example procedure, the `/u01/app/11.2.0/grid/` directory is the Grid Infrastructure Oracle Home.

   ```
   bash-3.00# /u01/app/11.2.0/grid/bin/crsctl add type sun.storage_proxy.type -basetype cluster_resource
   ```

3. Create **`sun.resource`** of type **`sun.storage_proxy.type`**.

   > **Note**
   > Ensure that all attribute values are enclosed in single quotes. Otherwise, the VALIDATE method of the `SUNW.scalable_asm_diskgroup_proxy` resource will fail the validation.

   ```
   bash-3.00# /u01/app/11.2.0/grid/bin/crsctl add res sun.scal-asmdg1-rs -type sun.storage_proxy.type \
   -attr "ACTION_SCRIPT='/opt/SUNWscor/dsconfig/bin/scproxy_crs_action' \
   HOSTING_MEMBERS='pfreak1 pfreak2' \
   CARDINALITY='2' \
   PLACEMENT='restricted' \
   ACL='owner:root:rwx,pgrp:oinstall:rwx,other::r--' \
   SCRIPT_TIMEOUT='20' \
   RESTART_ATTEMPTS='60'"
   ```

`sun.scal-asmdg1-rs`	The `SUNW.ScalDeviceGroup` resource name.
`-type sun.storage_proxy.type`	Specify the `sun.storage_proxy.type` resource type.
`ACTION_SCRIPT`	Specifies the `/opt/SUNWscor/dsconfig/bin/scproxy_crs_action` script.
`HOSTING_MEMBERS`	Specifies the Oracle Solaris Cluster resource group's node list entries that contain the `SUNW.ScalDeviceGroup` resource.
`CARDINALITY`	Sets the number of hosts that are defined in HOSTING_MEMBERS.
`PLACEMENT`	Set to `restricted`.

ACL	Sets the owner equal to `root` and the group equal to the ACL group entry for the Oracle ASM disk group. The following command displays the ACL group entry:
	`bash-3.00# /u01/app/11.2.0/grid/bin/crsctl stat res ora.DATA1.dg -p | grep ACL=`
	`ACL=owner:oragrid:rwx,pgrp:oinstall:rwx,other::r--`
	The example output shows that `oinstall` is the group entry.

`SCRIPT_TIMEOUT`	Set to `20`.
`RESTART_ATTEMPTS`	Set to `60`.

4. Verify that **sun.*resource*** is correctly defined.
   Output is similar to the following:

```
bash-3.00# /u01/app/11.2.0/grid/bin/crsctl stat res sun.scal-asmdg1-rs -p
NAME=sun.scal-asmdg1-rs
TYPE=sun.storage_proxy.type
ACL=owner:root:rwx,pgrp:oinstall:rwx,other::r--
ACTION_FAILURE_TEMPLATE=
ACTION_SCRIPT=/opt/SUNWscor/dsconfig/bin/scproxy_crs_action
ACTIVE_PLACEMENT=0
AGENT_FILENAME=%CRS_HOME%/bin/scriptagent
AUTO_START=restore
CARDINALITY=2
CHECK_INTERVAL=60
DEFAULT_TEMPLATE=
DEGREE=1
DESCRIPTION=
ENABLED=1
FAILOVER_DELAY=0
FAILURE_INTERVAL=0
FAILURE_THRESHOLD=0
HOSTING_MEMBERS=pfreak1 pfreak2
LOAD=1
LOGGING_LEVEL=1
NOT_RESTARTING_TEMPLATE=
OFFLINE_CHECK_INTERVAL=0
PLACEMENT=restricted
PROFILE_CHANGE_TEMPLATE=
RESTART_ATTEMPTS=60
SCRIPT_TIMEOUT=20
SERVER_POOLS=
START_DEPENDENCIES=
START_TIMEOUT=0
STATE_CHANGE_TEMPLATE=
STOP_DEPENDENCIES=
STOP_TIMEOUT=0
UPTIME_THRESHOLD=1h
bash-3.00#
```

5. Set the offline-restart dependency to add a dependency.
   a. Display the current dependency.

```
bash-3.00# /u01/app/11.2.0/grid/bin/crsctl stat res ora.DATA1.dg -p | grep
START_DEPENDENCIES
START_DEPENDENCIES=hard(ora.asm) pullup(ora.asm)
bash-3.00# clresource show -p resource_dependencies_offline_restart asm-data1-rs

=== Resources ===

Resource: asm-data1-rs
Resource_dependencies_offline_restart: asm-inst-rs

— Standard and extension properties —
```

b. Set the new dependency.

- If `asm-data1-rs` already exists, use the following command to set the dependency. Note that the command includes the plus (+) symbol:

```
bash-3.00# clresource set -p resource_dependencies_offline_restart+=scal-asmdg1-rs
asm-data1-rs
```

- If `asm-data1-rs` does not yet exist, use the following command to create the resource with the offline-restart dependency:

```
bash-3.00# clresource create -g asm-dg-rg \
-t SUNW.scalable_asm_diskgroup_proxy \
-p asm_diskgroups=data1 \
-p resource_dependencies_offline_restart=asm-inst-rs,scal-asmdg1-rs \
-d asm-data1-rs
```

c. Verify the configured dependency.

```
bash-3.00# /u01/app/11.2.0/grid/bin/crsctl stat res ora.DATA1.dg -p | grep
START_DEPENDENCIES
START_DEPENDENCIES=hard(ora.asm,sun.scal-asmdg1-rs) pullup(ora.asm)
bash-3.00# clresource show -p resource_dependencies_offline_restart asm-data1-rs
=== Resources ===

Resource: asm-data1-rs
Resource_dependencies_offline_restart: asm-inst-rs scal-asmdg1-rs

— Standard and extension properties —
```

Top

Removing an Oracle Grid Infrastructure Resource

This section contains the following procedures to remove an Oracle Grid Infrastructure resource.

Top

How to Remove a Dependency

This procedure shows how to set the offline-restart dependency to remove a dependency.

1. Display the current dependency.

```
bash-3.00# /u01/app/11.2.0/grid/bin/crsctl stat res ora.DATA1.dg -p | grep START_DEPENDENCIES
START_DEPENDENCIES=hard(ora.asm,sun.scal-asmdg1-rs) pullup(ora.asm)
bash-3.00# clresource show -p resource_dependencies_offline_restart asm-data1-rs
=== Resources ===

Resource: asm-data1-rs
Resource_dependencies_offline_restart: asm-inst-rs scal-asmdg1-rs

— Standard and extension properties —
```

2. Set the new dependency.
   Note that the command includes the minus (-) symbol.

```
bash-3.00# clresource set -p resource_dependencies_offline_restart-=scal-asmdg1-rs asm-data1-rs
```

3. Verify the modified dependency.

```
bash-3.00# /u01/app/11.2.0/grid/bin/crsctl stat res ora.DATA1.dg -p | grep START_DEPENDENCIES
START_DEPENDENCIES=hard(ora.asm) pullup(ora.asm)
bash-3.00# clresource show -p resource_dependencies_offline_restart asm-data1-rs

=== Resources ===

Resource: asm-data1-rs
Resource_dependencies_offline_restart: asm-inst-rs

— Standard and extension properties —
```

Top

How to Delete the **sun.*resource*** Resource

1. Become superuser.
2. Ensure that the dependency is removed, as described in "How to Remove a Dependency" , and that `sun.resource` is stopped.

```
bash-3.00# /u01/app/11.2.0/grid/bin/crsctl stop res sun.scal-asmdg1-rs
CRS-2500: Cannot stop resource 'sun.scal-asmdg1-rs' as it is not running
CRS-4000: Command Stop failed, or completed with errors.
```

3. Delete `sun.resource`.

```
bash-3.00# /u01/app/11.2.0/grid/bin/crsctl delete res sun.scal-asmdg1-rs
```

4. Verify that `sun.resource` is deleted.

```
bash-3.00# /u01/app/11.2.0/grid/bin/crsctl stat res sun.scal-asmdg1-rs -p
CRS-210: Could not find resource 'sun.scal-asmdg1-rs'.
```

Top

Network-Attached Storage Manual

This section discusses errors, omissions, and additions in the Oracle Solaris Cluster 3.3 With Network-Attached Storage Device Manual.

Extraneous Information for Oracle's Sun Storage 7000 Unified Storage System as a NAS Device

At the initial publication of this manual, part number 821-1556-10, Chapter 3 documented the use of Oracle's Sun Storage 7000 Unified Storage System as a NAS device with fencing enabled. Support for Oracle's Sun Storage 7000 Unified Storage System as a NAS Device with fencing has not been enabled in the 3.3 release. Chapter 3 has subsequently been removed from the manual and republished as part number 821-1556-11. If you obtained a copy of this manual before Chapter 3 was removed (821-1556-10), ignore all information in Chapter 3 and information elsewhere in the manual that refers to this functionality.

Top

Sun NAS and Network Appliance NAS Directories Must be Exported to All Cluster Nodes (6996565)

For Sun NAS devices, you must explicitly grant access to each node in the cluster. All nodes in the cluster must have access to configured directories. When you add the Sun NAS directories to a cluster, ensure that the Sun NAS device is properly configured and the directories the cluster will use have been exported to all cluster nodes.

For Network Appliance NAS devices, directories must be exported to all cluster nodes. When you add the Network Appliance NAS directories to a cluster, ensure that the Network Appliance NAS device is properly configured and the directories the cluster will use have been exported to all cluster nodes.

Top

Man Pages

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

rt_properties(5)

The `rt_properties` (5) man page lists incorrect API versions. The correct API versions for the following Sun Cluster or Oracle Solaris Cluster releases are listed below.

Sun Cluster 3.2 1/09 API version 9
Sun Cluster 3.2 11/09 API version 10
Oracle Solaris Cluster 3.3 API version 11

Top

clnasdevice(1CL)

Some sections refer to the use of Oracle's Sun Storage 7000 Unified Storage System as a NAS device, using the NAS device type `sun_uss`. Support for Oracle's Sun Storage 7000 Unified Storage System as a NAS Device has not been enabled in the 3.3 release. Ignore all information for the {{sun_uss} type.

Top

clressharedaddress(1CL)

The `clressharedaddress`(1CL) man page is missing from the `SUNWscman` package. Refer instead to the online version of the man page:

http://download.oracle.com/docs/cd/E19680-01/821-1263/clressharedaddress-1cl?l=en&a=view

Top

scds_hasp_check(3HA)

The following information is missing from the DESCRIPTION:

Resource dependencies are only checked within the same cluster context in which the function is executed, either global cluster

or zone cluster. Dependencies of the form *clustername*:*resourcename* (inter-cluster dependencies) are ignored. For example, if the only HAStoragePlus dependency is an inter-cluster dependency, the function returns the status code SCDS_HASP_NO_RESOURCE.

Top

SUNW.scalable_asm_diskgroup_proxy(5)

The SUNW.scalable_asm_diskgroup_proxy(5) is missing from the Oracle Solaris Cluster 3.3 release. The following are the man page contents.

Name

SUNW.scalable_asm_diskgroup_proxy, scalable_asm_diskgroup_proxy -- resource type implementation for the clustered Oracle Automated Storage Management (Oracle ASM) disk group managed by Oracle Solaris Cluster

Description

The SUNW.scalable_asm_diskgroup_proxy resource type represents the clustered Oracle ASM disk group in an Oracle Solaris Cluster configuration. This resource type is introduced in Oracle Solaris Cluster 3.3 for use starting with Oracle 11g release 2 configurations.

> **Note**
> The SUNW.scalable_asm_diskgroup_proxy resource type can only be used when using Oracle Grid Infrastructure for Clusters.

The SUNW.scalable_asm_diskgroup_proxy resource type is a multiple-master resource type. A single resource of this type can run on multiple nodes concurrently, but does not use network load balancing.

Each SUNW.scalable_asm_diskgroup_proxy resource represents a clustered Oracle ASM disk group. Each clustered Oracle ASM disk group is uniquely identified by the value of the asm_diskgroups extension property on the node where the instance is running. The Oracle ASM disk group resource should only be mounted if the Oracle ASM instance is available on the same cluster node. Additionally, Oracle ASM disk groups should only be mounted if any required storage resources are enabled on the cluster node. To ensure that these requirements are met, configure the Oracle ASM disk group resource as follows:

- Create a strong positive affinity between the Oracle ASM disk group resource group and the following resource groups:
  - The clustered Oracle ASM resource group
  - Any resource group that contains storage resources for Oracle files
- Create an offline-restart dependency between the clustered Oracle ASM disk group resource and the following resources:
  - The clustered Oracle ASM instance resource
  - Any storage resources for Oracle files that you are using

Create these dependencies and affinities when you configure clustered Oracle ASM disk group resources for the Oracle Solaris Cluster HA for Oracle data service or the Oracle Solaris Cluster Support for Oracle RAC data service. For more information about configuring resources for Oracle database instances, see Oracle Solaris Cluster Data Service for Oracle Guide or Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide.

To register this resource type and create instances of this resource type, use one of the following:

- Oracle Solaris Cluster Manager
- The clsetup(1CL) utility, specifying the option for configuring Oracle Solaris Cluster Support for Oracle Real Application Clusters
- The following sequence of Oracle Solaris Cluster maintenance commands:
  - To register this resource type, use the clresourcetype(1CL) command.
  - To create instances of this resource type, use the clresource(1CL) command.

Standard Properties

For a description of all standard resource properties, see the r_properties(5) man page.

Standard resource properties are overridden for this resource type as follows:

Prenet_start_timeout

Minimum	60
Default	300

Prenet_stop_timeout

Minimum	60
Default	300

Extension Properties

The extension properties of the `SUNW.scalable_asm_diskgroup_proxy` resource type are as follows.

asm_diskgroups

This property specifies the single-instance Oracle ASM disk group. If required, more than one single-instance Oracle ASM disk group can be specified as a comma separated list.

Data Type	String array
Range	Not applicable
Tunable	When disabled

debug_level

> **Note**
> All SQL*Plus and srvmgr messages that the Oracle ASM disk group resource issues are written to the log file `/var/opt/SUNWscor/oracle_asm/message_log.${RESOURCE}`.

This property indicates the level to which debug messages for the Oracle ASM disk-group resources are logged. When the debug level is increased, more debug messages are written to the system log `/var/adm/messages` as follows:

0	No debug messages
1	Function Begin and End messages
2	All debug messages and function Begin/End messages

Data Type	Integer
Range	0-2
Default	0
Tunable	Any time

proxy_probe_timeout

This property specifies the timeout value, in seconds, that the proxy monitor uses when checking the status of the clustered Oracle ASM disk group resource for which this resource is acting as a proxy.

Data Type	Integer
Range	5-120
Default	60
Tunable	Any time

proxy_probe_interval

This property specifies the interval, in seconds, between probes of the Oracle ASM disk group resource for which this resource is acting as a proxy.

Data Type	Integer
Range	5-120
Default	30
Tunable	Any time

Example - Creating a **scalable_asm_diskgroup_proxy** Resource

This example shows the commands for performing the following operations to create a scalable multi-master `scalable_asm_diskgroup_proxy` resource on a two-node cluster:

- Creating the `asm-dg-rg` resource group
- Registering the `SUNW.scalable_asm_diskgroup_proxy` resource type
- Setting the resource group affinity
- Adding the `asm-dg-rs` resource to the `asm-dg-rg` resource group
- Setting the `asm_diskgroups` extension property for one Oracle ASM disk group

The example makes the following assumptions:

- The `bash` shell is used.
- A resource group that is named `asm-inst-rg` exists and contains a resource of type `SUNW.scalable_asm_instance_proxy` that is named `asm-inst-rs`.
- A resource group that is named `scal-mp-rg` exists and contains a resource of type `SUNW.ScalMountPoint` that is named `scal-mp-rs`, for Oracle files.

```
phys-schost-1# clresourcetype register SUNW.scalable_asm_diskgroup_proxy

phys-schost-1# clresourcegroup create -S asm-dg-rg
phys-schost-1# clresourcegroup set -p rg_affinities=++asm-inst-rg asm-dg-rg

phys-schost-1# clresource create -g asm-dg-rg \
-t SUNW.scalable_asm_diskgroup_proxy \
-p asm_diskgroups=data1 \
-p resource_dependencies_offline_restart=asm-inst-rs,qfs-mp-rs \
-d asm-dg-rs

phys-schost-1# clresourcegroup online -M asm-dg-rg

phys-schost-1# clresource enable asm-dg-rs
```

Attributes

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWscor

See Also

Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide, Oracle Solaris Cluster Data Service for Oracle Guide, Oracle Solaris Cluster Data Services Planning and Administration Guide

Top

ORCL.obiee_cluster_controller(5)

The ORCL.obiee_cluster_controller(5) man page is missing from the Oracle Solaris Cluster 3.3 release. The following are the man page contents.

Name

ORCL.obiee_cluster_controller -- resource type implementation for HA-Oracle Business Intelligence Enterprise Edition (Oracle BI EE) Cluster Controller.

Description

The ORCL.obiee_cluster_controller resource type represents the HA Oracle BI EE Cluster Controller in an Oracle Solaris Cluster configuration.

The HA-Oracle BI EE Cluster Controller resource is configured in a failover resource group when creating an Oracle BI EE installation that consists of more than one BI EE Server. The node lists for the resource groups that contain the primary and secondary cluster controllers must not contain common nodes.

You must set the following properties for an Oracle BI EE Cluster Controller resource by using the clresource command.

Standard Properties

The standard resource property Failover is set for all failover resource types.

See the r_properties(5) man page for a complete description of the following resource properties.

Failover_mode

Default:	SOFT
Tunable:	Any time

Retry_count

Default:	2
Tunable:	Any time

Retry_interval

Default:	1330
Tunable:	Any time

Thorough_probe_interval

Default:	30
Tunable:	Any time

Extension Properties

BI_Install_Directory

Type string (required). This property is set to the absolute path of the Oracle BI software installation directory.

You can modify this parameter only when disabled.

BI_Data_Directory

Type string (required). This property is set to the absolute path of the Oracle BI software data directory.

You can modify this parameter only when disabled.

Run_64_bit

Type Boolean (required). Default is TRUE. If this property is set to TRUE, then the 64-bit version of the binary is run. If set to FALSE, then the 32-bit version of the binary is run.

You can modify this parameter only when disabled.

Cluster_Controller_Role

Type enumerated (required). Default is PRIMARY. This property determines whether the resource represents a primary or secondary cluster controller. Valid values are PRIMARY or SECONDARY.

You can modify this parameter only when disabled.

Debug_Level

Type enumerated (optional). Default is 0. Valid values are 0, 1, 2. This property specifies the debug level.
You can modify this parameter at any time.

See Also

`pmfadm`(1M), `scha_resource_get`(1HA), `clresourcetype`(1CL), `clresource`(1CL)

Oracle Solaris Cluster Data Services Planning and Administration Guide

Top

`ORCL.obiee_presentation_service`(5)

The `ORCL.obiee_presentation_service`(5) man page is missing from the Oracle Solaris Cluster 3.3 release. The following are the man page contents.

Name

`ORCL.obiee_presentation_service` -- resource type implementation for HA-Oracle Business Intelligence Enterprise Edition (Oracle BI EE) Presentation Service.

Description

The `ORCL.obiee_presentation_service` resource type represents the HA-Oracle BI EE Presentation Service in an Oracle Solaris Cluster configuration. The HA-Oracle BI EE Presentation Service resource can be configured in one of two ways:

- In a failover resource group when the BI EE Cluster Controller is not used.
- In a multi-master resource group when the BI EE Cluster Controller is used.

You must set the following properties for an Oracle BI EE Presentation Service resource by using the `clresource` command.

Standard Properties

The standard resource property `Failover` is set for all failover resource types.

See `r_properties`(5) for a complete description of the following resource properties.

Failover_mode

Default:	SOFT
Tunable:	Any time

Retry_count

Default:	2
Tunable:	Any time

Retry_interval

Default:	1330
Tunable:	Any time

Thorough_probe_interval

Default:	30
Tunable:	Any time

Extension Properties

BI_Install_Directory

Type string (required). This property is set to the absolute path of the Oracle BI software installation directory.

You can modify this parameter only when disabled.

BI_Data_Directory

Type string (required). This property is set to the absolute path of the Oracle BI software data directory.

You can modify this parameter only when disabled.

Run_64_bit

Type Boolean (required). Default is TRUE. If this property is set to TRUE, then the 64-bit version of the binary is run. If set to FALSE, then the 32-bit version of the binary is run.

You can modify this parameter only when disabled.

Debug_Level

Type enumerated (optional). Default is 0. Valid values are 0, 1, 2. This property specifies the debug level.

You can modify this parameter at any time.

See Also

`pmfadm`(1M), `scha_resource_get`(1HA), `clresourcetype`(1CL), `clresource`(1CL)

Oracle Solaris Cluster Data Services Planning and Administration Guide

Top

`ORCL.obiee_scheduler`(5)

The `ORCL.obiee_scheduler`(5) man page is missing from the Oracle Solaris Cluster 3.3 release. The following are the man page contents.

Name

`ORCL.obiee_scheduler` -- resource type implementation for HA Oracle BI EE Scheduler.

Description

The `ORCL.obiee_scheduler` resource type represents the HA-Oracle Business Intelligence Enterprise Edition (Oracle BI EE) Scheduler in an Oracle Solaris Cluster configuration. The HA-Oracle BI EE Scheduler resource can be configured in one of two ways:

- In a failover resource group when the BI EE Cluster Controller is not used.
- In a multi-master resource group when the BI EE Cluster Controller is used.

You must set the following properties for an Oracle BI EE Scheduler resource by using the `clresource` command.

Standard Properties

The standard resource property `Failover` is set for all failover resource types.

See `r_properties`(5) for a complete description of the following resource properties.

Failover_mode

Default:	SOFT
Tunable:	Any time

Retry_count

Default:	2
Tunable:	Any time

Retry_interval

Default:	1330
Tunable:	Any time

Thorough_probe_interval

Default:	30
Tunable:	Any time

Extension Properties

BI_Install_Directory

Type string (required). This property is set to the absolute path of the Oracle BI software installation directory.

You can modify this parameter only when disabled.

BI_Data_Directory

Type string (required). This property is set to the absolute path of the Oracle BI software data directory.

You can modify this parameter only when disabled.

Run_64_bit

Type Boolean (required). Default is TRUE. If this property is set to TRUE, then the 64-bit version of the binary is run. If set to FALSE, then the 32-bit version of the binary is run.
You can modify this parameter only when disabled.

Debug_Level

Type enumerated (optional). Default is 0. Valid values are 0, 1, 2. This property specifies the debug level.
You can modify this parameter at any time.

See Also

`pmfadm`(1M), `scha_resource_get`(1HA), `clresourcetype`(1CL), `clresource`(1CL)

Oracle Solaris Cluster Data Services Planning and Administration Guide

Top

`ORCL.obiee_server`(5)

The `ORCL.obiee_server`(5) man page is missing from the Oracle Solaris Cluster 3.3 release. The following are the man page contents.

Name

`ORCL.obiee_server` -- resource type implementation for HA Oracle BI EE server.

Description

The `ORCL.obiee_server` resource type represents the HA-Oracle Business Intelligence Enterprise Edition (Oracle BI EE) Server in an Oracle Solaris Cluster configuration. The HA-Oracle BI EE Server resource can be configured in one of two ways:

- In a failover resource group when the BI EE Cluster Controller is not used.
- In a multi-master resource group when the BI EE Cluster Controller is used.

You must set the following properties for an Oracle BI EE Server resource by using the `clresource` command.

Standard Properties

The standard resource property `Failover` is set for all failover resource types.

See `r_properties`(5) for a complete description of the following resource properties.

`Failover_mode`

Default:	SOFT
Tunable:	Any time

`Retry_count`

Default:	2
Tunable:	Any time

`Retry_interval`

Default:	1330
Tunable:	Any time

`Thorough_probe_interval`

Default:	30
Tunable:	Any time

Extension Properties

`BI_Install_Directory`

Type string (required). This property is set to the absolute path of the Oracle BI software installation directory.
You can modify this parameter only when disabled.

`BI_Data_Directory`

Type string (required). This property is set to the absolute path of the Oracle BI software data directory.
You can modify this parameter only when disabled.

Run_64_bit

Type Boolean (required). Default is TRUE. If this property is set to TRUE, then the 64-bit version of the binary is run. If set to FALSE, then the 32-bit version of the binary is run.
You can modify this parameter only when disabled.

BI_PROBE_USER

Type string (optional). This property contains the Business Intelligence (BI) user name to use for probing the health of the BI server.
You can modify this parameter only when disabled.

BI_PROBE_PASSWORD

Type string (optional). This property contains the Business Intelligence (BI) user password to use for probing the health of the BI server.
You can modify this parameter only when disabled.

Debug_Level

Type enumerated (optional). Default is 0. Valid values are 0, 1, 2. This property specifies the debug level.
You can modify this parameter at any time.

See Also

`pmfadm`(1M), `scha_resource_get`(1HA), `clresourcetype`(1CL), `clresource`(1CL)

Oracle Solaris Cluster Data Services Planning and Administration Guide

Top

`SUNW.oracle_server`(5)

Oracle Solaris Cluster 3.3 introduces a new extension property for the `SUNW.oracle_server` resource type. This extension property, `db_unique_name`, is required when a single-instance HA for Oracle data service is deployed using Oracle 11g release 2 Grid Infrastructure for Clusters. The `db_unique_name` extension property specifies the unique name of the single-instance database that is being deployed.

When the `db_unique_name` extension property is used, the following extension properties are ignored and can be omitted from the HA-Oracle configuration:

```
alert_log_file
connect_string
```

The following example shows the creation of the `ora-db1-rg` resource group, where the unique database name `DB1` is specified for the `db_unique_name` extension property for the database. Because the `db_unique_name` extension property is used, the `connect_string` and `alert_log_file` extension properties are omitted.

```
# clrs create -g ora-db1-rg \
-t SUNW.oracle_server \
-p oracle_home=/u01/app/ora11g/product/11.2.0/dbhome_1 \
-p oracle_sid=DB1 \
-p db_unique_name=DB1 \
-p resource_dependencies_offline_restart=asm-dg1-rs \
-d ora-db1-rs
```

Top

Online Help

This section discusses errors, omissions, and additions in the online help.

Oracle Solaris Cluster Manager

The following information is missing from the online help for Oracle Solaris Cluster Manager:

- The pages that contain the Resource Group Status Properties table and the Zone Cluster Status Properties table are missing the following fields:
    - Priority - The order in which resource groups are assigned to master nodes. A higher priority indicates a more important service. The default value is 500.
    - Preemption_Mode (Has_Cost/No_Cost/Never) - The likelihood that a resource group will be preempted from a node by a higher-priority resource group because of node overload. A Has_Cost value means that preempting this resource group has a cost associated with it. A No_Cost value indicates that the cost of preempting this resource group is zero. A Never value indicates that the resource group cannot be displaced from its current master to satisfy load limits.
    - Load Factors (Loadlimit@value) - How much of the load limit that the resource group consumes. The default value for each load factor is 0 and the maximum value is 1000.
- A new help page to describe how to create, edit, and delete a load limit:
    - Limitname - Name of the load limit.
    - Softlimit - The advisory upper boundary for a resource group load on a node or zone. The default value is 0, which indicates that no soft limit is imposed.
    - Hardlimit - Mandatory upper boundary for resource group load on a node or zone. The default value is null, which indicates that the "limit" is unlimited.

Top

What's New in Oracle Solaris Cluster 3.3

What's New in Oracle Solaris Cluster 3.3 Software

This section provides information related to new features, functionality, and supported products in Oracle Solaris Cluster 3.3 software.

The following new features and functionality are provided in patches to Oracle Solaris Cluster 3.3 software:

- NEW: Support for SWIFT Alliance and Gateway 7.0
- CLI Support to Configure Oracle RAC Instance Proxy Resource Group
- Support for Sun GlassFish Message Queue 4.4
- IP Address is Now Optional for Zone Cluster Nodes
- HA for Oracle Support for Third-Party Volume Managers with Clustered Oracle ASM Disk Groups
- Support for Apache Tomcat 5.5.28
- Oracle Solaris Cluster HA for Oracle TimesTen

The following new features and functionality are provided in the initial Oracle Solaris Cluster 3.3 release:

- NAS NFS File Systems Now Supported in a Zone Cluster
- Failover-File-System Monitoring
- Denying Cluster Services for a Non-Global Zone
- Trusted Extensions With Zone Clusters Support
- Oracle 11g Release 2 Support
- Load-Based Distribution of Resource Groups
- Node Rename
- Power Management User Interface
- Exporting Cluster File Systems to a Zone Cluster Using a Loopback Mount
- Wizard Updates for the `clsetup` Utility and Oracle Solaris Cluster Manager
- Wizard Support for Oracle ASM
- Support for Oracle WebLogic Server 10.3 in Failover and Clustered Configurations
- Enhanced Oracle E-Business Suite 3.2 Agent
- Oracle Solaris Cluster HA for MySQL Cluster
- Oracle Solaris Cluster HA for Oracle Business Intelligence Enterprise Edition
- Oracle Solaris Cluster HA for PeopleSoft Application Server

Top

CLI Support to Configure Oracle RAC Instance Proxy Resource Group

Support is added to the Oracle Solaris Cluster maintenance commands to configure an Oracle Real Application Clusters (Oracle RAC) instance proxy resource group when Oracle ASM is not also configured. This support is delivered starting in the following HA-Oracle patch versions:

- SPARC: 145335-01
- x86: 145336-01

For supporting configuration instructions, see Configuring an Oracle RAC Instance Proxy Resource Group Without ASM (CLI) .

Procedures for removing the Grid Infrastructure resource are also provided. See Removing an Oracle Grid Infrastructure Resource.

Top

Support for Sun GlassFish Message Queue 4.4

The installation restrictions can be removed when CR 6927826 is putback.

The following patches add support for Sun GlassFish Message Queue 4.4 to the HA for Java Message Queue data service:

- SPARC: 145337-02
- x86: 145338-01

After you install the appropriate patch, observe the following installation restrictions for Message Queue 4.4:

- Do not use the GlassFish Message Queue installer.
- Install only the SVR4 packages (not the IPS packages) by using the `pkgadd` command.

Otherwise, you can following procedures as published in Oracle Solaris Cluster Data Service for Sun Java System Message Queue Guide.

Top

IP Address is Now Optional for Zone Cluster Nodes (6961705)

Specifying an IP address and network interface card (NIC) for each zone cluster node is now optional. This new functionality is included in the Oracle Solaris Cluster 3.3 145333-02 Patch, and it is not documented in the Oracle Solaris Cluster 3.3 `clzonecluster`(1CL) man page. If you do not configure an IP address for each zone cluster node, two things will occur:

1. That specific zone cluster will not be able to configure NAS devices for use in the zone cluster. The cluster uses the IP address of the zone cluster node when communicating with the NAS device, so not having an IP address prevents cluster support for fencing NAS devices.
2. The cluster software will activate any Logical Host IP address on any NIC.

HA for Oracle Support for Third-Party Volume Managers with Clustered Oracle ASM Disk Groups

Support is added to the HA for Oracle data service for third-party volume managers configured with clustered Oracle ASM disk groups. This support is delivered in the following HA for Oracle patches:

- SPARC: 145335-01
- x86: 145336-01

For supporting configuration instructions, see Support for Third-Party Volume Managers with Clustered Oracle ASM Disk Groups .

Support for Apache Tomcat 5.5.28

Support is added to the HA for Apache Tomcat data service for Apache Tomcat versions 5.5.28 and greater. This support requires an HA for Apache Tomcat patch and a change in how the data service is configured. See the following information for details:

- `SUNWsctomcat` Resource Probe Fails with Tomcat Version 5.5.28 & 5.5.29 (6964640)
- Data Service for Apache Tomcat Guide

Top

Oracle Solaris Cluster HA for Oracle TimesTen

Oracle Solaris Cluster now supports a data service for Oracle TimesTen In-Memory Database. For information about installing and configuring HA for Oracle TimesTen, see the Oracle Solaris Cluster Data Service for Oracle TimesTen In-Memory Database Guide .

> **Note**
> The following patch is required for this data service:
>
> - SPARC: 145339-01
> - x86: 145340-01

Top

NAS NFS File Systems Now Supported in a Zone Cluster

Network-Attached Storage (NAS) NFS file systems can be used as storage devices for applications that run in a zone cluster with fencing support. You can enable this feature with the Oracle Solaris Cluster Manager GUI or by using the `clnasdevice add -Z` command. For instructions, see Oracle Solaris Cluster Network-Attached Storage Device Manual.

Top

Failover-File-System Monitoring

Oracle Solaris Cluster 3.3 actively monitors mounted file systems and detects whether a mounted file system is truly available to an application, meaning the underlying storage is accessible. When necessary, the cluster performs a corrective failover. See Chapter 2 of the Oracle Solaris Cluster Software Data Services Planning and Administration Guide.

Top

Denying Cluster Services for a Non-Global Zone

You can turn off cluster functionality for a selected non-global zone, so that a root user logged into one of these zones will not be able to discover or disrupt operation of the cluster. For instructions, see How to Deny Cluster Services For a Non-Global Zone and How to Allow Cluster Services for a Non-Global Zone in the Oracle Solaris Cluster Data Services Planning and Administration Guide.

Top

Trusted Extensions With Zone Clusters Support

Oracle Solaris Cluster on Solaris Trusted Extensions is now supported, but is limited to certain hardware and software configurations. Contact your Oracle representative for the latest information about the currently supported configurations.

For guidelines and configuration procedures, see Oracle Solaris Cluster Software Installation Guide.

Top

Oracle 11g Release 2 Support

Oracle RAC 11g release 2 is supported with Oracle Solaris Cluster. The following restrictions apply to certain Oracle 11g release 2 data-service configurations:

- HA-Oracle and Oracle RAC - The use of Oracle ASM disk groups is supported only on hardware RAID. The use of Oracle ASM disk groups on devices that are managed by a volume manager is not supported.
- HA-Oracle and Oracle RAC - To configure the Oracle ASM resource groups and their resources, use only CLI commands. Do not use Oracle Solaris Cluster Manager or the `clsetup` utility to configure an Oracle ASM resource group.
- Oracle RAC - To configure an Oracle RAC instance proxy resource group, the configuration tools to use are as follows:
  - If the Oracle RAC database does not use Oracle ASM, you must use either the Oracle Solaris Cluster Manager GUI or the `clsetup` utility.
  - If the Oracle RAC database does use Oracle ASM, you must use only the Oracle Solaris Cluster CLI.

Guidelines and procedures in Oracle Solaris Cluster Data Service for Oracle Guide and Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide are valid for Oracle 11g release 2 unless otherwise stated in the documentation.

Top

Load-Based Distribution of Resource Groups

Load-based resource group distribution lets you create a workload distribution policy. This policy considers system capacity and loading while starting or failing over resources. For configuration procedures, see Configuring the Distribution of Resource Group Load Across Nodes in the Oracle Solaris Cluster Software Installation Guide] and How to Configure Load Limits on a Node in the Oracle Solaris Cluster System Administration Guide.

Top

Node Rename

You can use the `clnode rename` command to change the name of a node that is part of an Oracle Solaris Cluster configuration. You must rename the Oracle Solaris hostname before you can rename the node. See How to Rename a Node in the Oracle Solaris Cluster System Administration Guide.

Top

Power Management User Interface

The Oracle Solaris power management user interface can now turn Oracle Solaris Cluster components and systems on and off in order to save energy. See Overview of Administering the Cluster in the Oracle Solaris Cluster System Administration Guide.

Top

Exporting Cluster File Systems to a Zone Cluster Using a Loopback Mount

You can make a cluster file system available to a zone cluster by mounting the file system in the global cluster and then performing a loopback mount in the zone-cluster node. Oracle Solaris Cluster can manage loopback mounts for cluster file systems on UFS and on Veritas File System. See How to Add a Local File System to a Zone Cluster in the Oracle Solaris Cluster Software Installation Guide] and How to Set Up the HAStoragePlus Resource for Cluster File Systems in the Oracle Solaris Cluster Data Services Planning and Administration Guide.

Top

Wizard Updates for the **`clsetup`** Utility and Oracle Solaris Cluster Manager

The `clsetup` utility and Oracle Solaris Cluster Manager wizards were expanded to support configuring a `SUNW.vucmm_framework` resource group and resources for an Oracle Real Application Clusters configuration. Information about using `SUNW.vucmm_framework`, also called the multiple-owner volume-manager framework, is in Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide.

Top

Wizard Support for Oracle ASM

You can now configure Oracle ASM for HA-Oracle and Oracle RAC by using wizards in the `clsetup` utility and the Oracle Solaris Cluster Manager GUI. You can also configure HA-Oracle and Oracle RAC in a zone cluster. Information is in Oracle Solaris Cluster Data Service for Oracle Guide and Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide.

Top

Support for Oracle WebLogic Server 10.3 in Failover and Clustered Configurations

Oracle Solaris Cluster now supports Oracle WebLogic Server 10.3 in the failover and clustered ("multi-master") modes of operation. This includes support for Oracle WebLogic Server 10.3 servers in zone clusters. Information is located in Oracle Solaris Cluster Data Service for WebLogic Server Guide.

Top

Enhanced Oracle E-Business Suite 3.2 Agent

The Oracle E-Business Suite 3.2 agent now supports Parallel Concurrent Processing. Information is located in Oracle Solaris Cluster Data Service for Oracle E-Business Suite Guide.

Top

Oracle Solaris Cluster HA for MySQL Cluster

Oracle Solaris Cluster now supports a data service for MySQL Cluster. For information about installing and configuring HA for MySQL Cluster, see the Oracle Solaris Cluster Data Service for MySQL Cluster Guide. Additional information to support MySQL Cluster has been added to the existing Oracle Solaris Cluster Data Service for MySQL Guide.

Top

Oracle Solaris Cluster HA for Oracle Business Intelligence Enterprise Edition

Oracle Solaris Cluster now supports a data service for Oracle Business Intelligence Enterprise Edition. For information about installing and configuring HA for Oracle Business Intelligence Enterprise Edition, see the Oracle Solaris Cluster Data Service for Oracle Business Intelligence Enterprise Edition Guide.

Top

Oracle Solaris Cluster HA for PeopleSoft Application Server

Oracle Solaris Cluster now supports a data service for the PeopleSoft application server on SPARC based platforms. For information about installing and configuring HA for PeopleSoft application server, see the Oracle Solaris Cluster Data Service for PeopleSoft Enterprise Guide .

Top

Support for SWIFT Alliance Access and Gateway 7.0

SWIFT Alliance Access and Gateway 7.0 Now Supported on Oracle Solaris Cluster 3.3

The following products are now supported with Oracle Solaris Cluster 3.3:

- SWIFT Alliance Access 7.0
- SWIFT Alliance Gateway 7.0

Oracle Solaris Cluster HA for Alliance Access 7.0

The Oracle Solaris Cluster 3.3 HA agent now works with SWIFT Alliance Access version 7.0. Alliance Access is a trademark of SWIFT. This section contains Configuration Restrictions and instructions for Installing and Configuring version 7.0. Patch 146241-02 is required for SWIFT Alliance Access version 7.0.

Configuration Restrictions

This section provides a list of software and hardware configuration restrictions that apply to Oracle Solaris Cluster HA for Alliance Access only.

Caution - Your data service configuration might not be supported if you do not observe these restrictions.

- You can configure the Oracle Solaris Cluster HA for Alliance Access only as a HA agent and not as a scalable agent.
- You can install the Alliance Access software on a global file system. Best practice is to use a failover file system.
- Only one Alliance Access instance is supported by this agent.
- You can install Alliance Access 7.0 only with the embedded database option.
- You can upgrade to Alliance Access 7.0 only with the non-root installation procedure.

Installing and Configuring Alliance Access 7.0

1. Create the resources for Alliance Access.
    a. Create a resource group for Alliance Access.

```
# clresourcegroup create [-n node-zone-list] swift-rg
```

`-n node-zone-list`
Specifies a comma-separated, ordered list of zones that can master the resource group. The format of each entry in the list is node. In this format, node specifies the node name and zone specifies the name of a non-global Oracle Solaris zone. To specify the global zone, or to specify a node without non-global zones, specify only node. This list is optional. If you omit this list, the global zone of each cluster node can master the resource group.

b. Create a logical host.
Add the hostname and IP address in the `/etc/inet/hosts` file on all cluster nodes or zones that can master the resource group. Register the logical host and add it to the resource group.

```
# clreslogicalhostname create -g swift-rg -l swift-lh swift-saa-lh-rs
```

c. Create the device group and file system.
d. Create an HAStoragePlus resource.
Create an HAStoragePlus failover resource to contain the Alliance Access application and configuration data instead of using the global file system.

```
# clresource create -g swift-rg
-t SUNW.HAStoragePlus
-x FilesystemMountPoints=/global/saadg/alliance swift-ds
```

e. Bring the resource group online.

```
# clresouregroup online -M swift-rg
```

f. Create the configuration directory.
This directory contains Alliance Access information and creates a link from the `/usr`

```
# cd /global/saadg/alliance
# mkdir swa
# ln -s /global/saadg/alliance/swa /usr/swa
```

NOTE
If you install Alliance Access in a sparse root zone, that is if the `/usr` directory is inherited in read-only mode through a loopback mount, the link needs to be created within the global zone.

g. For Alliance Access 7.0:

```
# mkdir -p /global/saadg/alliance/var/opt/swift
```

On all cluster nodes or zones that can master the resource group:

```
# ln -s /global/saadg/alliance/var/opt/swift /var/opt/swift
```

2. Install Alliance Access 7.0 software.
Perform the following steps on all cluster nodes or zones that can master the resource group.
a. Create the user `all_adm` and the groups `alliance` and `sagnlg` on all cluster nodes or zones that can master the resource group with the same user ID and group ID. Also, create a project called `swift` and assign the users `all_adm` to it.

```
# groupadd -g groupid alliance
# groupadd -g groupid sagsnlg
# useradd -m -g alliance -G sagsnlg -d /export/home/all_adm -s /usr/bin/ksh all_adm
# projadd -U all_adm swift
```

b. Set the values of the resource controls for the project `swift`.

```
# chown -R all_adm:alliance /global/saadg/alliance
# projmod -s -K "project.max-sem-ids=(privileged,1320,deny)" swift
# projmod -s -K "project.max-shm-ids=(privileged,1500,deny)" swift
# projmod -s -K "project.max-shm-memory=(privileged,4294967295,deny)" swift
# projmod -s -K "project.max-msg-ids=(privileged,800,deny)" swift
# projmod -s -K "process.max-sem-nsems=(privileged,512,deny)" swift
# projmod -s -K "process.max-sem-ops=(privileged,512,deny)" swift
# projmod -s -K "process.max-msg-qbytess=(privileged,10485760,deny)" swift
# projmod -s -K "process.max-msg-messages=(privileged,8192,deny)" swift
# projmod -s -K "process.max-stack-size=(basic,33554432,deny)" swift
# projmod -s -K "process.max-data-size=(basic,8.0EB,deny)" swift
# projmod -s -K "process.max-file-descriptor=(basic,4096,deny)" swift
```

The previous values are examples. For more accurate values, refer to the latest SWIFT documentation Release Notes.

c. Assign the project `swift` as the default project for `all_adm` by editing the file `/etc/user_attr` and typing the following line at the end of the file:

```
all_adm::::project=swift
```

3. Add a symbolic link and entries.
   a. Add the symbolic link `/usr/swa` on all cluster nodes or zones that can master the resource group. See Step 1f.
   b. Add the symbolic link `/var/opt/swift` on all cluster nodes or zones that can master the resource group. See Step 1g.
   c. As the root user, run the `saa_rootpostinstall.ksh` script on all cluster nodes or zones that can master the resource group:

```
# /global/saadg/alliance/install/saa_rootpostinstall.ksh
```

4. Install Alliance Access Remote API (RA).
   Install RA after Alliance Access on shared storage using the following options:
   Instance `RA1` (default), user `all_adm`

Oracle Solaris Cluster HA for Alliance Gateway 7.0

The Oracle Solaris Cluster 3.3 HA agent now works with SWIFT Alliance Gateway version 7.0. Alliance Gateway is a trademark of SWIFT. In this version, SWIFTAlliance Gateway is renamed as Alliance Gateway. You can use the Oracle Solaris Cluster HA for Alliance Gateway with both the product names. HA for Alliance Gateway 7.0 is supported for Oracle Solaris Cluster 3.3 only, and not for Oracle Solaris Cluster releases prior to 3.3.

This section contains Configuration Restrictions, instructions for Installing and Configuring, and instructions for Debugging Alliance Gateway 7.0.

Configuration Restrictions

The configuration restrictions in this section apply only to Oracle Solaris Cluster HA for Alliance Gateway.

Caution - Your data service configuration might not be supported if you do not observe these restrictions.

- Oracle Solaris Cluster HA for Alliance Gateway supports Alliance Gateway version 6.0, 6.1, 6.3, and 7.0.

- You can install the SWIFTNet Link and the Alliance Gateway software only on a failover file system. If WebSphere MQ client software is needed for the operation of Alliance Gateway, install WebSphere MQ client software on the local file system in the same path on each cluster node or zone that can master the resource group. The default path is `/opt/mqm`.
- You cannot configure Oracle Solaris Cluster HA for Alliance Gateway as a scalable data service.
- If the Alliance Gateway instance is started through the cluster resource, do not stop it manually. Instead, stop Alliance Gateway by disabling the corresponding cluster resource.

> **NOTE**
> The Oracle Solaris Cluster HA for Alliance Gateway can be configured to run in a whole root or sparse root non-global zone for Oracle Solaris Cluster HA for Alliance Gateway version 6.0, 6.1, 6.3, or 7.0 if required.

Installing and Configuring Alliance Gateway 7.0

To enable Oracle Solaris Cluster HA for Alliance Gateway to make Alliance Gateway highly available, additional installation and configuration operations are required. These operations supplement the standard installation and standard configuration of Alliance Gateway.

During a standard installation, Alliance Gateway is installed with a physical hostname. To enable Alliance Gateway to run in a cluster, you must modify Alliance Gateway to use a logical hostname.

For information about the standard installation and standard configuration of Alliance Gateway, see the appropriate documentation available in the Alliance Gateway CD-ROM.

To perform this procedure, you need the server root directory (the path to the application binaries). You can install the binaries on the local disks or on the cluster file system. For a discussion of the advantages and disadvantages of each location, see Chapter 1 in Oracle Solaris Cluster Data Services Planning and Administration Guide.

1. Create a resource group for Alliance Gateway.

```
# clresourcegroup create [-n node-zone-list] sag-rg
```

   `-n node-zone-list`
   Specifies a comma-separated, ordered list of zones that can master the resource group. The format of each entry in the list is node. In this format, node specifies the node name and zone specifies the name of a non-global Oracle Solaris zone. To specify the global zone, or to specify a node without non-global zones, specify only node. This list is optional. If you omit this list, the global zone of each cluster node can master the resource group.

2. Create a logical host.
   A logical host is required before you install Alliance Gateway.
   a. Add the hostname and the IP address in the `/etc/inet/hosts` file on all cluster nodes or zones that can master the resource group.
   b. Register the logical host, and then add the logical host resource to the resource group.

```
# clreslogicalhostname -g sag-rg -h swiftgatewayhost sag-lh-rs
```

3. Create the device group and file systems.
   Although you can use a global file system, create an HAStoragePlus failover resource to contain the Alliance Gateway application and configuration data.
   - To create a device group and a file system for Alliance Gateway, see Planning the Global Devices, Device Groups, and Cluster File Systems in the Oracle Solaris Cluster Software Installation Guide.
   - To create an HAStoragePlus failover resource, see Enabling Highly Available Local File Systems in the Oracle Solaris Cluster Data Services Planning and Administration Guide.
     This procedure uses `/swift` as the path that contains the Alliance Gateway application and configuration data.

```
# clresource create -g sag-rg -t SUNW.HAStoragePlus -x filesystemMountPoints=/swift sag-ds
```

4. Bring the resource group online to enable the IP address and access to the storage.

```
# clresourcegroup online -M sag-rg
```

5. Create the following directories and symbolic links before the installation.
   Without these symbolic links, the Alliance Gateway application will not be installed in the correct location. You must install the Alliance Gateway application in the correct location to ensure failover capabilities.
   Perform this procedure on the node where the resource group for Alliance Gateway 7.0 is online:

   ```
   # mkdir -p /swift/etc/opt/swnet /swift/etc/opt/secrets_backup
   # chown swnet:swnetg /swift/etc/opt/swnet /swift/etc/opt/secrets_backup
   # chmod -R 0555 /swift/etc
   # chmod 0775 /swift/etc/opt/swnet /swift/etc/opt/secrets_backup
   # mkdir -p /swift/var/opt/swnet /swift/var/opt/swift
   # chown swnet:swnetg /swift/var/opt/swnet /swift/var/opt/swift
   # chmod -R 0555 /swift/var
   # chmod 0775 /swift/var/opt/swnet /swift/var/opt/swift
   # mkdir -p /swift/home/swnet
   # chown swnet:swnetg /swift/home/swnet
   ```

   Perform this procedure on all cluster nodes:

   ```
   # ln -s /swift/home/swnet /export//home/swnet
   # ln -s /swift/etc/opt/swnet /etc/opt/swnet
   # ln -s /swift/etc/opt/secrets_backup /etc/opt/secrets_backup
   # ln -s /swift/var/opt/swnet /var/opt/swnet
   # ln -s /swift/var/opt/swift /var/opt/swift
   ```

6. Install the WebSphere MQ client packages, if required. WebSphere MQ client software is software that guarantees and load-balances connections between the gateway and remote SWIFTNet Link systems. If you chose this type of Alliance Gateway installation and have the appropriate license, install the WebSphere MQ client packages. The WebSphere MQ client package must be installed locally by using the same installation directory on all cluster nodes or zones that can master the resource group.
7. Install SWIFTNet Link.
   You must install SWIFTNet Link on the node or zone where the resource group for Alliance Gateway is online. You must also install SWIFTNet Link in the directory on the file system managed by the HAStoragePlus resource that you created in Step 3.
   a. Follow the instructions in your SWIFTNet Link documentation. To refer to the SWIFTNet Link documentation, obtain the SWIFTNet Link CD-ROM.
   b. Specify the directory on which to install the failover data service: `/swift/snl`.
   c. Install any patches for SWIFTNet Link, if required.
   d. The Hardware Security Module (HSM) Installation/Configuration wizard starts up automatically when you click Finish at the end of SWIFTNet Link installation. Click Cancel to exit the HSM Installation/Configuration wizard.
   e. Verify that connectivity with SWIFTNet Switch is established by running the SWIFTNet Link selftest command. For more information, see the SWIFTNet Admin Services: Operational Interface on the Alliance Gateway CD-ROM.
   f. Stop the SWIFTNet Link instance by running the SWIFTNet Link stop command.
8. Configure and register HSM on the primary node as follows.
   a. Log on to the primary node or zone as the SWIFTNet Link instance owner with the `swnet` account.
   b. Run the following command:

      ```
      # perl SwHSMDiskClone.pl -a SETUP
      ```

   c. Log off.
   d. Log on as the SWIFTNet Link instance owner with the swnet account.
   e. Use the HSM administration tool SwHSM to configure and register the HSM boxes connected to the primary node. See the SWIFTNet Link Installation and Administration Guide for UNIX for details.

> **NOTE**
> During the configuration and registration of the HSM, you must enter the boot IP address of the primary node in the SWIFTNet Link host IP address field of the Register screen. This source IP address is used to establish the connection. Do not use the logical host address that is used by the service.

9. Register HSM on the secondary node or zone.
   a. Switch the resource group for Alliance Gateway to the secondary node or zone.

      ```
      # clresourcegroup switch -n node2 sag-rg
      ```

   b. Log on to the secondary node as the SWIFTNet Link instance owner with the `swnet` account.
   c. Run the following command:

      ```
      # perl SwHSMDiskClone.pl -a SETUP
      ```

   d. Log off.
   e. Log on as the SWIFTNet Link instance owner with the swnet account.
   f. Use the HSM administration tool SwHSM to register the HSM boxes connected to the secondary node or zone. See the SWIFTNet Link Installation and Administration Guide for UNIX for details.

> **NOTE**
> During the registration of the HSM, you must enter the boot IP address of the secondary node or zone in the SWIFTNet Link host IP address field of the Register screen. This source IP address is used to establish the connection. Do not use the logical host address that is used by the service.

> **NOTE**
> You must be aware that the only difference between Step 8 and this step is that you configure and register the HSM on the first node or zone in Step 8, whereas you only register the HSM on the secondary node or zone in this step.

10. Install Alliance Gateway software.
    You must install Alliance Gateway on the node or zone where the resource group for Alliance Gateway is online. You must also install Alliance Gateway in a directory on the file system managed by the HAStoragePlus resource that you created in Step 3. In this procedure, `/swift/SWIFTAlliance/Gateway` is being used as the directory. You might install the Alliance Gateway patches, if necessary.
    - Follow the instructions in your Alliance Gateway documentation. To refer to the Alliance Gateway documentation, obtain the Alliance Gateway CD-ROM.
    - Use the logical IP address as the IP with which the Alliance Gateway software communicates with remote hosts.
11. Synchronize all nodes or zones with installation-specific changes to user files and system files.
    Add all entries in `/etc/system` and `/etc/services` added by the installation on the first node to the secondary node or zone that can master the resource group. The entries in /etc/system must be applied to the global zone.

Debugging Oracle Solaris Cluster HA for Alliance Gateway 7.0

Oracle Solaris Cluster HA for Alliance Gateway has an option file in `/opt/SUNWscsag/etc` directory that enables you to set a debug flag.

Normally, the debug information for Oracle Solaris Cluster software does not reside in the log files. You must edit the `syslog.conf` file to enable logging of messages of level debug. You can edit the `syslog.conf` to log those messages in another file.

1. Edit the `/etc/syslog.conf` file.
   a. Change `daemon.notice` to `daemon.debug`.
      The following output from the `grep daemon /etc/syslog.conf` command shows that `daemon.debug` has not been set.

```
grep daemon /etc/syslog.conf
*.err;kern.debug;daemon.notice;mail.crit
/var/adm/messages
*.alert;kern.err;daemon.err
operator
```

b. Restart `syslogd`.

```
# pkill -1 syslogd
```

The `syslogd` command forces `syslog` to reread its configuration file and account for changes.
The following output from the `grep daemon /etc/syslog.conf` command shows that `daemon.debug` has been set.

```
grep daemon /etc/syslog.conf
*.err;kern.debug;daemon.debug;mail.crit
/var/adm/messages
*.alert;kern.err;daemon.err
operator
```

2. Edit the `/opt/SUNWscsag/etc/config` file.
   Change the line DEBUG= to DEBUG=ALL or DEBUG=`resource`.

How to Debug a Failed Restart for the Alliance Gateway Resource

If the Alliance Gateway instance failed to restart on a cluster node or zone, but the instance did successfully fail over to another cluster node or zone, perform the following steps on the node or zone where the restart failed.

1. Log in as the Alliance Gateway UNIX user swnet.
2. Use the `ipcs`(1) command to determine if there are active inter-process communication (IPC) facilities for the swnet user.

```
$ ipcs -opqs | grep swnet
```

3. If you see output for the command you ran in Step 2, verify that there is no Alliance Gateway instance running on this node or zone.
   To clean up the IPC resource from the swnet user, either use the `ipcrm`(1) command for each facility for the swnet user, or reboot the node or zone.
   You can switch the resource group with the Alliance Gateway instance back to this node or zone only if there are no IPC resources left for the swnet user.

Features Nearing End of Life

Features Nearing End of Life

There are no features nearing end of life.

Top

Product Name Changes

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.

Top

Product Name Changes

This section provides information about product name changes for applications that Oracle Solaris Cluster software supports. Depending on the Oracle Solaris Cluster software release that you are running, your Oracle Solaris Cluster documentation might not reflect the following product name changes.

Top

Current Product Name	Former Product Name
Oracle Solaris Cluster	Sun Cluster (also Solaris Cluster and Java Availability Suite)
Oracle Solaris Cluster Geographic Edition	Sun Cluster Geographic Edition
Oracle Solaris Cluster data services	Sun Cluster data services
Oracle Solaris Cluster Data Service for Oracle Grid Engine	Sun Cluster Data Service for Sun Grid Engine
Oracle Solaris Cluster Data Service for Oracle VM Server for SPARC	Sun Cluster Data Service for LDoms Guest Domain
Oracle Solaris Cluster Manager	Sun Cluster Manager

Top

Supported Products

Supported Products

This section describes the supported software and memory requirements for Oracle Solaris Cluster 3.3 software.

Top

Data Services

Contact your Oracle sales representative for the complete list of supported data services (agents) and application versions. Links to the documentation for many data service agents are available here.

Top

File Systems

Solaris 10 SPARC

File System	Additional Information
Solaris UFS	
Solaris ZFS	Not supported for the `/globaldevices` file system
Sun StorEdge QFS	

Sun QFS 5.0 and 5.1 - Standalone file system	Supported Data Services: All failover data services External Volume Management: Solaris Volume Manager only
QFS 4.6 - Standalone file system	Supported Data Services: All failover data services External Volume Management: Solaris Volume Manager, VxVM
QFS 4.6, 5.0, and 5.1 - Shared QFS file system	Supported Data Services: Oracle RAC External Volume Management: Solaris Volume Manager for Sun Cluster
QFS 4.6, 5.0, and 5.1 - Shared QFS clients outside the cluster (SC-COTC)	Supported Data Services: None; only a shared file system is supported External Volume Management: No external volume manager is supported
QFS 4.6, 5.0, and 5.1 - HA-SAM Failover	Supported Data Services None; only a shared file system is supported External Volume Management: No external volume manager is supported
Veritas File System components that are delivered as part of Veritas Storage Foundation 5.0 and 5.1. Version 5.0 requires a minimum of MP3 RP3.	

Top

Solaris 10 x86

File System	Additional Information
Solaris UFS	
Solaris ZFS	Not supported for the `/globaldevices` file system
Sun StorEdge QFS	
Sun QFS 5.0 and 5.1 - Standalone file system	Supported Data Services: All failover data services External Volume Management: Solaris Volume Manager only
QFS 4.6 - Standalone file system	Supported Data Services: All failover data services External Volume Management: Solaris Volume Manager, VxVM
QFS 4.6, 5.0, and 5.1 - Shared QFS file system	Supported Data Services: Oracle RAC External Volume Management: Solaris Volume Manager for Sun Cluster
QFS 4.6, 5.0, and 5.1 - Shared QFS clients outside the cluster (SC-COTC)	Supported Data Services: None; only a shared file system is supported External Volume Management: No external volume manager is supported
QFS 4.6, 5.0, and 5.1 - HA-SAM Failover	Supported Data Services None; only a shared file system is supported *External Volume Management:* No external volume manager is supported
Veritas File System components that are delivered as part of Veritas Storage Foundation 5.0 and 5.1. Version 5.0 requires a minimum of MP3 RP3.	

Top

Memory Requirements

Oracle Solaris Cluster 3.3 software requires the following memory requirements for every cluster node:

- Minimum of 1 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.

Top

Solaris Operating System (OS)

Oracle Solaris Cluster 3.3 software and Quorum Server software requires one of the following versions of the Solaris OS:

- Solaris 10 - Solaris 10 10/09, Oracle Solaris 10 9/10

> Note -
> Oracle Solaris Cluster 3.3 software does not support multiple versions of Solaris software in the same running cluster.

Top

Sun Logical Domains (LDoms)

This Oracle Solaris Cluster release supports Sun Logical Domains 1.2 and 1.3 software. For version 1.2, a minimum of patch 142840-05 is required.

Top

Sun Management Center

This Oracle Solaris Cluster release supports Sun Management Center software versions 3.6.1 and 4.0.

Top

Sun StorageTek Availability Suite

This Oracle Solaris Cluster release supports Sun StorageTek Availability Suite 4.0 software. Support requires a minimum of patch 123246-05 for SPARC or 123247-05 for x86.

Top

Volume Managers

This Oracle Solaris Cluster release supports the following volume managers.

Solaris 10 SPARC

Volume Manager	Cluster Feature
Solaris Volume Manager	Solaris Volume Manager for Sun Cluster
Veritas Volume Manager (VxVM) components that are delivered as part of Veritas Storage Foundation 5.0 and 5.1. Version 5.0 requires a minimum of MP3 RP3.	VxVM 5.0 and VxVM 5.1 cluster feature (with RAC only)

Top

Solaris 10 x86

Volume Manager	Cluster Feature
Solaris Volume Manager	Solaris Volume Manager for Sun Cluster
Veritas Volume Manager components that are delivered as part of Veritas Storage Foundation 5.0 and 5.1. Version 5.0 requires a minimum of MP3 RP3. Veritas Volume Replicator (VVR) and Fast Mirror Resynchronization (FMR) are not supported on the x86 platform.	Not applicable - Oracle Solaris Cluster 3.3 software does not support the VxVM cluster feature on the x86 platform.

Top

Oracle Solaris Cluster Security Hardening

Oracle Solaris Cluster Security Hardening uses the Solaris Operating System hardening techniques recommended by the Sun BluePrints™ program to achieve basic security hardening for clusters. The Solaris Security Toolkit automates the implementation of Oracle Solaris Cluster Security Hardening.

The Oracle Solaris Cluster Security Hardening documentation is available http://blogs.sun.com/security/entry/reference_security_blueprints . From this URL, scroll down to the 2003 heading to locate the article "Securing the Sun Cluster 3.x Software." The documentation describes how to secure Sun Cluster 3.x deployments in a Solaris environment. The description includes the use of the Solaris Security Toolkit and other best-practice security techniques recommended by Oracle security experts. The following data services are supported by Oracle Solaris Cluster Security Hardening:

- Oracle Solaris Cluster HA for Apache
- Oracle Solaris Cluster HA for Apache Tomcat
- Oracle Solaris Cluster HA for DHCP
- Oracle Solaris Cluster HA for DNS
- Oracle Solaris Cluster HA for MySQL
- Oracle Solaris Cluster HA for NFS
- Oracle Solaris Cluster HA for Oracle
- Oracle Solaris Cluster HA for Oracle E-Business Suite
- Oracle Solaris Cluster HA for Oracle Grid Engine
- Oracle Solaris Cluster Support for Oracle Real Application Clusters
- Oracle Solaris Cluster HA for PostgreSQL
- Oracle Solaris Cluster HA for Samba
- Oracle Solaris Cluster HA for Siebel
- Oracle Solaris Cluster HA for Solaris Containers
- Oracle Solaris Cluster HA for SWIFTAlliance Access
- Oracle Solaris Cluster HA for SWIFTAlliance Gateway
- Oracle Solaris Cluster HA for Sun Java System Directory Server
- Oracle Solaris Cluster HA for Sun Java System Message Queue
- Oracle Solaris Cluster HA for Sun Java System Messaging Server
- Oracle Solaris Cluster HA for Sun Java System Web Server
- Oracle Solaris Cluster HA for Sybase ASE
- Oracle Solaris Cluster HA for WebLogic Server
- Oracle Solaris Cluster HA for WebSphere MQ
- Oracle Solaris Cluster HA for WebSphere MQ Integrator

Top

Compatibility Issues

Compatibility Issues

This section contains information about Oracle Solaris Cluster compatibility issues with other products.

- Additional Oracle Solaris Cluster framework compatibility issues are documented in "Planning the Oracle Solaris Cluster Configuration" in Oracle Solaris Cluster Software Installation Guide.
- Additional Oracle Solaris Cluster upgrade compatibility issues are documented in "Upgrade Requirements and Software

Support Guidelines" in Oracle Solaris Cluster Upgrade Guide.
- For other known problems or restrictions, see Known Issues and Bugs.

Top

Oracle Database 11.2.0.2

For Oracle Database version 11.2.0.2, when using the Oracle Universal Installer (OUI) to install Oracle Grid Infrastructure, you must specify the Oracle ASM discovery string so that only devices that are accessible by the Oracle Grid Infrastructure user will match. If other devices are matched that the Oracle Grid Infrastructure user cannot access, the installation fails.

For example, instead of setting the disk discovery path as `/dev/did/rdsk/d*`, list only the specific disks that are accessible by the Oracle Grid Infrastructure software owner, such as `/dev/did/rdsk/d1s5,/dev/did/rdsk/d2s5`.

For more information, see https://support.us.oracle.com/oip/faces/secure/km/DocumentDisplay.jspx?id=1303209.1.

Top

Undefined Symbol Errors After **`scinstall`** on First Boot on x86

On the Solaris 10 10/09 OS, after running `scinstall` and booting into an x86 cluster for the first time, the node fails to come up and displays multiple "undefined symbol" messages, such as the following:

```
Sep 13 09:42:58 phys-node-1 genunix: [ID 819705 kern.notice]
/kernel/misc/amd64/cl_haci: undefined symbol
Sep 13 09:42:58 phys-node-1 genunix: [ID 826211 kern.notice]  '_memcpy'
Sep 13 09:42:58 phys-node-1 genunix: [ID 472681 kern.notice] WARNING: mod_load: cannot
load module 'cl_haci'
Sep 13 09:42:58 phys-node-1 genunix: [ID 819705 kern.notice]
/kernel/misc/amd64/cl_haci: undefined symbol
...
```

This issue is caused by Solaris bug 6396272, which is fixed in Solaris patch 142910-17. To avoid this problem, follow procedures in "How to Apply a Rebooting Patch (Node)" in the Oracle Solaris Cluster System Administration Guide to apply this patch to each cluster node.

Top

Solaris Upgrade to Solaris 10 9/10 Overwrites the **`hostname.`*`adp`*** Contents (6971566)

When the Solaris OS is upgraded to the Oracle Solaris 10 9/10 release, the upgrade overwrites the `/etc/hostname.`*`adp`* files. To ensure that you do not lose your configuration data, perform the following steps before and after the upgrade:

1. Before you upgrade the OS, make a backup of the contents of each `/etc/hostname.`*`adp`* file.
2. After the upgrade but before you reboot the system, copy the backup contents of each `/etc/hostname.`*`adp`* file to the current `/etc/hostname.`*`adp`* files.
3. After you reboot, if there was IPMP group information in the `/etc/hostname.`*`adp`* files, verify that IPMP groups are created correctly. Use either the `cluster` command or `ifconfig` commands.

If you already upgraded the OS to Oracle Solaris 10 9/10 and your `/etc/hostname.`*`adp`* files were overwritten, follow procedures in Solaris administration documentation to recreate the files, as well as affected IPMP groups.

Top

Zone Root Path Gets Modified on Doing a Live Upgrade with ZFS Root on **`cluster`** Branded Zone (6852390)

When a cluster that uses a ZFS root file system and which has zone cluster configured is upgraded by using Live Upgrade, the zone root path is modified such that the upgraded zone cannot boot.

To avoid this problem, apply the following patch before you begin live upgrade of the cluster:

- SPARC: 121430-45
- x86: 121431-46

Top

Cluster Zone Won't Boot Up After Live Upgrade on ZFS Root (6955669)

For a global cluster that uses ZFS for the root file system and which has zone clusters configured, when using Live Upgrade to upgrade to Solaris 10 8/10, the upgraded boot environment does not boot.

Contact your Oracle support representative to learn whether a patch or workaround is available.

Top

Webconsole Branding Must Be Changed to Match Oracle Brand (6925641)

The Oracle Solaris Cluster Manager GUI depends on changes to Java Web Console that are not contained in the Solaris 10 11/09 release. The necessary changes are available in the following required patches:

- 125952-20 (SPARC)
- 125953-20 (x86)

Top

The Command **zoneadmd** Should Mount **cluster** Branded Zones As **native** in Alternate Roots: Solaris Upgrade Fails in Presence of Zone Clusters (6874636)

Problem Summary: If a zone cluster is configured on a global-cluster node that runs the Solaris 10 OS, the use of JumpStart to upgrade the OS of that global-cluster node fails with a warning message like the following:

> Warning
> The Solaris Version (Solaris 10) on slice `c1t1d0s0` cannot be upgraded.
>
> A non-global zone could not be mounted.

The problem affects all JumpStart upgrades of the Solaris 10 OS to the Solaris 10 10/09 OS, if there are zone clusters installed on the global-cluster node. For example, an upgrade from the Solaris 10 5/09 OS to the Solaris 10 10/09 OS fails.

> Note
> (1) The problem does not occur if no zone clusters are installed on the cluster being upgraded.
> (2) The problem does not occur if there are zone clusters, but they exist only in the `Configured` state.

Workaround: Add two files to the Solaris installation miniroot image on the JumpStart server, as shown in the steps below. In these instructions, server is the name of the JumpStart server that you use to install the upgrade. Perform all steps as superuser.

Perform Step 1 from one node of the cluster that is to be upgraded.

1. Copy the zone-cluster `config.xml` and `platform.xml` files from a node that is to be upgraded to a location on the JumpStart server.

```
# cp /usr/lib/brand/cluster/config.xml.upgrade /net/server/some_dir/config.xml
# cp /usr/lib/brand/cluster/platform.xml /net/server/some_dir/platform.xml
```

Perform Steps 2-6 on the JumpStart server.

2. Create a new Solaris installation image.

```
# cd path_to_Solaris_install_image/Solaris_10/Tools
# ./setup_install_server path_to_new_Solaris_install_image
# cd
```

3. Unpack the miniroot.

```
# /boot/solaris/bin/root_archive unpackmedia path_to_new_Solaris_install_image
path_to_miniroot
```

4. Put the zone-cluster brand files in the miniroot.

```
# mkdir path_to_miniroot/usr/lib/brand/cluster
# cp /net/server/some_dir/config.xml path_to_miniroot/usr/lib/brand/cluster
# cp /net/server/some_dir/platform.xml path_to_miniroot/usr/lib/brand/cluster
```

5. Pack the miniroot and place it in the new installation image.

```
# /boot/solaris/bin/root_archive packmedia path_to_new_Solaris_install_image
path_to_miniroot
```

6. Run the `addclient` command from the new installation image location for each cluster node to be upgraded.

```
# cd path_to_new_Solaris_install_image/Solaris_10/Tools
# ./addclient your_configuration_details
```

Perform this workaround just once for each JumpStart server and for each Solaris update to which you plan to upgrade.

Top

Accessibility Features for People With Disabilities

To obtain accessibility features that have been released since the publishing of this media, consult Section 508 product assessments that are available from Oracle on request to determine which versions are best suited for deploying accessible solutions.

Top

Solaris Volume Manager GUI

The Enhanced Storage module of Solaris Management Console (Solaris Volume Manager) is not compatible with Oracle Solaris Cluster software. Use the command-line interface or Oracle Solaris Cluster utilities to configure Solaris Volume Manager software.

Top

Patches and Required Firmware Levels

Patches and Required Firmware Levels

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

- Applying an Oracle Solaris Cluster 3.3 Core Patch
- Removing an Oracle Solaris Cluster 3.3 Core Patch
- Patch Management Tools
- Patch for Cluster Support for Sun StorageTek 2530 Array
- My Oracle Support

If you are upgrading to Oracle Solaris Cluster 3.3 software, see the Oracle Solaris Cluster Upgrade Guide. Applying an Oracle Solaris Cluster 3.3 Core patch does not provide the same result as upgrading the software to the Oracle Solaris Cluster 3.3 release.

Note
Read the patch README file before applying or removing any patch.

You must be a registered My Oracle Support user to view and download the required patches 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 My Oracle Support .

Top

Applying an Oracle Solaris Cluster 3.3 Core Patch

Complete the following procedure to apply an Oracle Solaris Cluster 3.3 core patch. Ensure that all nodes of the cluster are maintained at the same patch level.

How to Apply an Oracle Solaris Cluster 3.3 Core Patch

Caution
If an Oracle Solaris Cluster 3.3 core patch is removed, any resources that were upgraded in Step 3 must be downgraded to the earlier resource type versions. The procedure for downgrading requires planned downtime of these services. Therefore, do not perform Step 3 until you are ready to commit the Oracle Solaris Cluster 3.3 core patch permanently to your cluster.

1. Install the patch by using the usual rebooting patch procedure for a core patch.
2. Verify that the patch has been installed correctly on all nodes and is functioning properly.
3. Perform Resource Type upgrades to any new versions of the resource type available. Run `clsetup` to obtain the list of new resource types. For more details, see the documentation on "Upgrading a Resource Type" in the Oracle Solaris Cluster Data Services Planning and Administration Guide.

For information about registering a resource type, see "Registering a Resource Type" in the Oracle Solaris Cluster Data Services Planning and Administration Guide.

Top

Removing an Oracle Solaris Cluster 3.3 Core Patch

Complete the following procedure to remove an Oracle Solaris Cluster 3.3 core patch.

How to Remove an Oracle Solaris Cluster 3.3 Core Patch

1. List the resource types on the cluster.

```
# clresourcetype list
```

2. If you have upgraded to any new resource types after applying the core patch, follow the directions in "How to Remove a Resource Type" in the Oracle Solaris Cluster Data Services Planning and Administration Guide.
3. Remove the Oracle Solaris Cluster 3.3 core patch from each node on which you installed the patch.

```
# patchrm patch-id
```

4. Reboot into cluster mode all of the nodes from which you removed the Oracle Solaris Cluster 3.3 core patch. Rebooting all of the nodes from which you removed the Oracle Solaris Cluster 3.3 core patch before rebooting any unaffected nodes ensures that the cluster is formed with the correct configuration information on all nodes. If all nodes on the cluster were patched with the core patch, you can reboot the nodes into cluster mode in any order.
5. Reboot any remaining nodes into cluster mode.

For instructions about rebooting nodes into cluster mode, see "How to Reboot a Cluster Node" in the Oracle Solaris Cluster System Administration Guide.

Top

Patch Management Tools

Information about patch management options for the Solaris OS is available at the web site for Oracle Enterprise Manager Ops Center (formerly Sun Ops Center).

The following tools are part of the Solaris OS. Refer to the version of the manual that is published for the Solaris OS release that is installed on your system:

- Information for using the Solaris patch management utility, `patchadd`, is provided in Solaris Administration Guide: Basic Administration at http://download.oracle.com/docs/cd/E19253-01/817-1985/index.html.
- Information for using Solaris Live Upgrade to apply patches is provided in the Solaris installation guide for Live Upgrade and upgrade planning at http://download.oracle.com/docs/cd/E19253-01/821-0438/index.html.

If some patches must be applied when the node is in noncluster mode, you can apply them in a rolling fashion, one node at a time, unless a patch's instructions require that you shut down the entire cluster. Follow procedures in How to Apply a Rebooting Patch (Node) in Oracle Solaris Cluster System Administration Guide to prepare the node and boot it into noncluster mode. For ease of installation, consider applying all patches at once to a node that you place in noncluster mode.

Top

Patch for Cluster Support for Sun StorageTek 2530 Array

The Sun StorageTek Common Array Manager (CAM) software, minimum required Version 6.0.1, provides SCSI3 or PGR support for the Sun StorageTek 2530 array for up to three nodes. The patch is not a required upgrade for the Sun StorEdge 6130, 2540, 6140, and 6540, and StorageTek FLX240, FLX280 and FLX380 platforms. The CAM patch is available from the Sun Download Center.

Top

My Oracle Support

The My Oracle Support Online Web site provides 24-hour access to the most up-to-date information regarding patches, software, and firmware for Oracle products. Access the site at My Oracle Support for the most current matrixes of supported software, firmware, and patch revisions.

Before you install Oracle Solaris Cluster 3.3 software and apply patches to a cluster component (Solaris OS, Oracle Solaris Cluster software, volume manager software, data services software, or disk hardware), review each README file that accompanies the patches that you retrieved. All cluster nodes must have the same patch level for proper cluster operation.

For specific patch procedures and tips on administering patches, see "Patching Oracle Solaris Cluster Software and Firmware" in Oracle Solaris Cluster System Administration Guide.

Top

Known Bugs in Oracle Solaris Cluster 3.3

Known Issues and Bugs

The following known issues and bugs affect the operation of the Oracle Solaris Cluster 3.3 release. Bugs and issues are grouped into the following categories:

- Administration
- Data Services
- GUI
- Installation
- Localization
- Upgrade

Also see Compatibility Issues and Documentation Issues .

Top

Administration

Oracle's SPARC T3-4 Fails During Reboot (6993321)

Problem Summary: During a reboot, Oracle's SPARC T3-4 server with four processors fails to connect to the Oracle Solaris Cluster

framework. For more detailed information, see 6986241. Error messages similar to the following appear:

```
Sep 20 15:18:53 svc.startd [8]: svc:/system/pools:default: Method or service exit timed out. Killing contract 29.
Sep 20 15:18:53 svc.startd [8]: svc:/system/pools:default: Method "/lib/svc/method/svc-pools start" failed due to
signal KILL.
...
Sep 20 15:20:55 solta svc.startd [8]: system/pools:default failed: transitioned to maintenance (see 'svcs -xv' for
details)
...
Sep 20 15:22:12 solta INITGCHB: Given up waiting for rgmd.
...
Sep 20 15:23:12 solta Cluster.GCHB_resd: GCHB system error: scha_cluster_open failed with 18
Sep 20 15:23:12 solta : No such process
```

Workaround: Use the `svccfg` command to increase the service timeout to 300 seconds. Boot into non-cluster mode and perform the following commands:

```
svccfg -s svc:/system/pools setprop start/timeout_seconds = 300
svcadm refresh svc:/system/pools
```

After you perform these commands, boot into cluster mode.
Top

Default Quorum Timeout of 25 Seconds May Not Be Adequate for Remote Quorums on IP Networks With Routing Involved (6974930)

Problem Summary: Cluster reconfiguration has a timeout of 25 seconds for the completion of quorum operations. This timeout value is deemed appropriate for typical quorum-device connectivity technologies, such as Fiber Channel or Direct-Attached SCSI. But a 25-second timeout might not be adequate for other technologies that are based on IP networks, especially with routing in such networks. This caveat also applies to any SCSI device, regardless of topology, that might be experiencing I/O timeouts and retrys, even if the latter eventually do succeed.

If you see failures of cluster nodes with the message `CMM: Unable to acquire the quorum device`, the cluster reconfiguration's quorum operations did not complete within the default timeout of 25 seconds. This might be a problem with the quorum device or with the path to the device. Check that both the quorum device and the path to it are functional.

Workaround: If the problem continues and you want to keep using the same quorum configuration, perform the following steps to increase the quorum operations timeout on all nodes of the cluster. The example below shows how to increase the timeout.

1. Become superuser.
2. On each cluster node, edit the `/etc/system` file as superuser to set the timeout to a high value.
   The following example sets the timeout to 700 seconds.

   ```
   phys-schost# vi /etc/system
   ...
   set cl_haci:qd_acquisition_timer=700
   ```

3. From one node, shut down the cluster.

   ```
   phys-schost-1# cluster shutdown -g0 -y
   ```

4. Boot each node back into the cluster.
   Changes to the `/etc/system` file are initialized after the reboot.

> **NOTE**
> For Oracle Real Application Clusters (Oracle RAC), do not change the default quorum timeout of 25 seconds. In certain split-brain scenarios, a longer timeout period might lead to the failure of Oracle RAC VIP failover, due to the VIP resource timing out. If the quorum device being used is not conforming with the default 25 seconds timeout, use a different quorum device.

Top

The `cluster status -t node` Command Does Not Work (6973987)

Problem Summary: The `cluster status -t node` command does not work as expected and gives the following error message:

```
% cluster status -t node
cluster: (C103603) Object type "node" does not have status.
```

Workaround: Run the following command instead to query the comprehensive node status:

```
% clnode status [-v] [-l] [-m] [-Z zonecluster | global | all ]  [+ | node ...]
```

Top

Newly Added Zone-Cluster Node Is Not Recognized by `clresourcegroup` or by Some `clzonecluster` Options (6957622)

Problem Summary: This issue occurs for a zone cluster that is configured on a subset of nodes in the global cluster and that has at least one zone-cluster node in the Online state. If a new zone-cluster node is added to that zone cluster, operations related to resource groups and resources do not recognize the newly added node as a valid node of the zone cluster. For example, an attempt to create a resource group with a node list that contains the newly added node results in an `Invalid node` error. Similarly, if you run the command `scha_cluster_get -O ALL_ZONES -Z` *`zonecluster`* from the global zone, the output does not report the newly added zone-cluster node.

Workaround: After a new node is added to a zone cluster, you must reboot all global-cluster nodes before they will recognize the new zone-cluster node.

To avoid this issue, ensure that all zone-cluster nodes that you might need are created at the time you create the zone cluster itself.

Top

When Using ZFS on Replicated Devices, Device Group Can Be Moved Away, Leaving Application Unable to Write to Storage (6848624)

Problem Summary: In a campus cluster configuration using Hitachi Universal Replicator and ZFS, if a user issues a manual switchover of a device group containing replicated ZFS storage pools (zpools), the operation results in the following undesirable state:

- The zpools do not migrate to the new primary node and will have their state changed to write-disabled.
- The application resource groups using the zpools do not migrate to the new primary node.
- The associated raw-disk device groups migrate to the new primary node and have their state set to read-write access.

Workaround: Do not manually switch over a device group that contains zpools that are replicated by Hitachi Universal Replicator.

Solaris Cluster Manager Fails to Come Up in a 16-Node Cluster (6594485)

Problem Summary: For a 16-node cluster, the Oracle Solaris Cluster Manager GUI is not usable.

Workaround: Use instead the `clsetup` utility or the Oracle Solaris Cluster administrative commands.

Top

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.

Top

Load Balancing Weight Is Not Updated if the Scalable Resource Is Disabled (6883814)

Problem Summary: If the load balancing weight of a scalable resource is changed when the resource is in a disabled state, the changed weight is not reflected in the client request distribution after the resource is enabled.

Workaround: Execute the command again after the resource is enabled.

Top

Deleting a Resource Which Has Dependent `clrs: (C979102) "<rs>" is not present in the property list` (6970087)

Public Summary: Unable to delete a resource which has dependencies that were set by using the `-F` option in the `clresourcegroup delete` command.

Workaround: Unset the dependencies before you delete the resources.

Top

The `global_fencing` Property Code is Broken When the Value is Changed to `prefer3` (6879360)

Problem Summary: Oracle Solaris Cluster attempts to verify that a storage device fully supports SCSI-3 PGR before allowing the user to set its fencing property to `prefer3`. This verification might succeed when it should fail.

Workaround: Ensure that a storage device is certified by Oracle Solaris Cluster for use with SCSI-3 PGR before changing the fencing setting to `prefer3`.

Top

Global-Cluster Node Cannot Get Zone Cluster Resource Group or Resource Information if its Zone-Cluster Node Has Never Been Booted Into Cluster Mode (6770391)

Problem Summary: This problem occurs for a zone cluster that is configured with resources and resource groups. If a zone-cluster node has not booted into cluster mode even once, then running the `clresourcegroup status -Z` *`zonecluster`* and `clresource status -Z` *`zonecluster`* commands in the global zone returns an `unexpected error` message. Similarly, running the `scha_resourcegroup_get -Z` *`zonecluster`* and `scha_resource_get -Z` *`zonecluster`* commands in the global zone returns error 18, if the zone-cluster nodes are not yet booted into cluster mode. On such nodes, Oracle Solaris Cluster Manager does not report the status of resource groups and resources for the zone cluster.

Workaround: Boot the zone-cluster nodes into cluster mode by running the following command in the global zone, where *`nodelist`* is the comma-separated list of nodes on which the zone-cluster nodes are not yet booted into cluster mode:

```
# clzonecluster boot -n nodelist zonecluster
```

This command starts the cluster services that are required to work around this problem. You only need to boot the zone-cluster nodes once when the physical node is booted. After that, you can halt or boot the zone-cluster nodes at your discretion, for as long as the physical node remains up.

Top

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.

Top

More Validation Checks Needed When Combining DIDs (6605101)

Problem Summary: The `scdidadm` and `cldevice` commands are unable to verify that replicated SRDF devices 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 DID device instances specified are replicas of each other and that they belong to the specified replication group.

Top

Data Services

HAStoragePlus Resource Configured for Zone Cluster with VxFS File System Fails to Come Online (6972831)

Problem Summary: When a VxFS file system is configured with an HAStoragePlus resource for a zone cluster, the resource fails to come online because of a file-system check failure. The error message is similar to the following:

```
File system check of mountpoint (rawdevice) failed: (39)
rawdevice: BADSUPERBLOCK AT BLOCK 16: MAGIC NUMBER WRONG
rawdevice: USE AN ALTERNATE SUPERBLOCK TO SUPPLY NEEDED INFORMATION;
rawdevice: e.g. fsck [-F ufs] -o b=# [special ...]
rawdevice: where # is the alternate super block. SEE fsck_ufs(1M).
rawdevice: UNEXPECTED INCONSISTENCY; RUN fsck MANUALLY.
Failed to fsck : mountpoint
```

Workaround: Ensure that the VxFS file system information is kept in the `/etc/vfstab` file of the global zone along and that the file system is added to the zone cluster configuration by using the `clzonecluster` command.

Top

11gR2 CRS May Not Start if **`crs_framework`** Resource Was Disabled During Boot Up (6971060)

Problem Summary: This issue applies to cluster nodes that have Oracle 11g release 2 installed and are using the RAC framework.

If a node is rebooted with the RAC framework disabled, the node upon booting attempts to start the Oracle Grid Infrastructure. This results in the cluster node being in a state where the Grid Infrastructure cannot be started or stopped until the node is rebooted, either with the RAC framework enabled or the Oracle Grid Infrastructure autostart disabled.

Workaround: Before you reboot a node with the RAC framework disabled, first disable the Oracle Grid Infrastructure autostart. See the Oracle 11g release 2 documentation for the command to use to disable autostart of the Grid Infrastructure.

Top

The Oracle WebLogic Server Proxy Server Resource Fails to Start (6970737)

Problem Summary: The Oracle WebLogic Server Proxy Server resource fails to start if the Managed Servers are offline.

Workaround: The WebLogic Server Proxy Server resource within a WebLogic Server proxy resource group should be dependent on WebLogic Server Managed Server resources. Strong resource dependencies ensure that the WebLogic Server Proxy Server is brought online before Managed Server resources are brought online. Use this command to specify the dependencies between the WebLogic Server proxy resource and the Managed Server resource:

```
# clresource create \
-t SUNW.wls \
-g wls-proxy-rg \
-p resource_dependencies=wls-mgd1-rs,wls-mgd2-rs ... \
wls-proxy-rs

# clresource set \
-p resource_dependencies=wls-mgd1-rs,wls-mgd2-rs ... \
wls-proxy-rs
```

Top

The Oracle WebLogic Server Stops Core Dumps (6970466)

Problem Summary: The Oracle WebLogic Server stops core dumps if the `smooth_shutdown` property is set to `true` and the `server_url` property does not include the port number.

Workaround: The WLS `server_url` property value should be in the form of `http://hostname:port`.

_
Top

Cannot Create RAC Proxy Resource When User Has the Same Group ID for Grid and Database Homes (6969233)

Problem Summary: The `SUNW.scalable_rac_server_proxy` resource cannot be created if the software installations of both the Oracle 11g release 2 Database and Grid Infrastructure used the same DBA group ID.

Workaround: Reinstall the Oracle 11g release 2 Database software using a different DBA group ID than that which was used for the Grid Infrastructure installation.

Top

The **`SUNW.crs_framework`** Resource Does Not Attempt a Restart When Oracle Clusterware Goes Offline (6966668)

Problem Summary: If Oracle Clusterware software fails on a node, the `SUNW.crs_framework` resource type does not automatically restart the software. This problem occurs whenever Oracle Clusterware software itself goes down and is unable to restart itself.

Workaround: Restart the Oracle Clusterware software manually. See the appropriate Oracle release documentation to determine the appropriate commands.

Top

`SUNWsctomcat` Resource Probe Fails with Tomcat Version 5.5.28 & 5.5.29 (6964640)

Problem Summary: Starting with version 5.5.28, the Apache Tomcat server changed the behavior of the `http get` command.

Workaround: Due to this change, an alternative probe algorithm is supplied in the following Oracle Solaris Cluster patches for HA-Apache Tomcat:

- SPARC: 145341-01
- x86: 145342-01

This probe algorithm uses the `wget` command, bypassing the normal probe. An HA for Apache Tomcat data service using Apache Tomcat version 5.5.28 or greater must use the `wget` probe algorithm. You can use the `wget` probe algorithm for versions earlier than 5.5.28 as well.

In addition, supporting documentation for this functionality is added to or replaces what is published in Oracle Solaris Cluster Data Service for Apache Tomcat Guide. See Data Service for Apache Tomcat Guide for details.

Top

HAStoragePlus Resource Configured in Scalable Resource Group with Cluster File System Stays at "Starting" State Indefinitely (6960386)

Problem Summary: When the `/etc/vfstab` file entry for a cluster file system has a mount-at-boot value of "no" and the cluster file system is configured in a `SUNW.HAStoragePlus` resource that belongs to a scalable resource group, the `SUNW.HAStoragePlus` resource fails to come online. The resource stays in the Starting state until `prenet_start_method` is timed out.

Workaround: In the `/etc/vfstab` file's entry for the cluster file system, set the mount-at-boot value to "yes".

Top

Configuration of Resource Groups With **`start`*`app_name`*** Script Fails (6941251)

Problem Summary: The `start`*`app_name`* script fails to create resource groups, because an incorrect `nodelist` is passed to the `clresourcegroup create` command.

Workaround: You must specify `-n nodelist`, even for the default node list.

Top

Scalable Resource Fails To Start Due To Failure of Unrelated IPMP Group (6938555)

Problem Summary: A scalable resource that depends on a `SUNW.SharedAddress` resource fails to come online due to a failure of an IPMP group that is on a subnet not used by the shared-address resource. The following messages can be 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 scalable resource that failed to start before.

Alternatively, remove the failed IPMP group (see the `ifconfig`(1M) man page) and restart the scalable resource that failed to start before.

> **NOTE**
> Network interfaces in the failed IPMP group will no longer be monitored by Oracle Solaris Cluster after the IPMP group is removed. Resources that depend on the failed IPMP group can no longer be created or started until the IPMP group is recreated and repaired.

Top

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 running in different zone clusters.

Workaround: The scalable applications should not be configured to bind to `INADDR_ANY` as the local IP address or they should bind to a port that does not conflict with another scalable application.

Top

Validation Should Fail for a Scalable Resource With the **`Outgoing_connection`** Property Set and With Multiple Entry Points (6886818)

Problem Summary: The problem occurs when the `Outgoing_Connection` property is set on a scalable resource and the resource is configured with multiple failover resource groups, each containing one or more scalable addresses. Such a configuration is not caught during validation, but TCP connections for the scalable service might be disrupted as a result.

Workaround: If the `Outgoing_Connection` property is enabled for a scalable resource, include all its scalable addresses in a single failover resource group.

Top

GUI

Data Service Configuration Wizards for Apache, NFS, HA-Oracle, and SAP Do Not Set an Offline-Restart Dependency on an HAStoragePlus Resource (6947273)

Problem Summary: The data-service configuration wizards in the Oracle Solaris Cluster GUI or the `clsetup` utility do not set an offline-restart dependency on an HAStoragePlus resource from the application. This is a regression in the wizards that was introduced in this release with the new failover-file-system monitoring feature.

If the resource is an existing resource that was previously created by the wizard, it will continue to operate without failover-file-system monitoring until the user explicitly upgrades the HAStoragePlus resource type.

Workaround: To enable failover-file-system monitoring for an existing resource, , take offline the HAStoragePlus resource, upgrade it to the 3.3 version resource type, and modify its resource dependencies appropriately.

To create a new resource with failover-file-system monitoring enabled, do not use the data-service wizards to create the HAStoragePlus resources where file-system monitoring is to be enabled. Use the command-line interface instead.

Top

RAC Wizard Configures QFS MDS Resources for Different Zone Clusters Under Same QFS MDS Resource Group by Default (6887179)

Problem Summary: The RAC wizard, by default, creates the QFS MDS resources for different zone clusters and the global cluster under the same QFS MDS resource group. The zone cluster nodes can be halted or rebooted independently of each other. So if the resources are in the same resource group, the administration of one resource could inadvertently impact another.

Workaround: The wizard provides an option to edit the auto-generated resource and resource group names in the review panel. Use this option to edit the QFS MDS resource group name, which is auto-generated by the wizard, and give it a value that is different from the one that already exists on the cluster.

Top

Installation

Autodiscovery Does Not Work for the **`qlge`** Driver for PCIe FCoE CNA (6939847)

Problem Summary: During an Oracle Solaris Cluster installation, auto discovery for the `qlge` driver for the PCIe LP and PCIe ExpressModule FCoE Converged Network Adapters (CNAs) does not work. The following products are affected:

- Oracle's Sun Storage 10GbE FCoE PCIe CNA
- Oracle's Sun Storage 10GbE FCoE ExpressModule CNA

Workaround: When you run the `scinstall` utility and you are prompted for the interconnect adapters, select Other and type the name of each `qlge` interface.

Top

Localization

Result of System Requirements Checking Is Wrong (6495984)

Problem Summary: When you use the `installer` utility in the Simplified Chinese and Traditional Chinese locales to install Oracle Solaris Cluster software, the software that checks the system requirements incorrectly reports that the swap space is 0 Mbytes.

Workaround: Ignore this reported information. In these locales, you can run the following command to determine the correct swap space:

```
# df -h | grep swap
```

Top

Upgrade

The Command **`zoneadmd`** Should Mount **`cluster`** Branded Zones As **`native`** in Alternate Roots: Solaris Upgrade Fails in Presence of Zone Clusters (6874636)

Problem Summary: This problem affects all JumpStart upgrades of the Oracle Solaris 10 OS to at least the Solaris 10 10/08 OS, if there are zone clusters installed on the global-cluster node. If a zone cluster is configured on a global-cluster node that runs the Oracle Solaris 10 OS, the use of JumpStart to upgrade the OS of that global-cluster node fails with a warning message similar to the following:

```
Warning
The Oracle Solaris Version (Solaris 10) on slice c1t1d0s0 cannot be upgraded.

A non-global zone could not be mounted.
```

For example, an upgrade from the Solaris 10 5/09 OS to the Solaris 10 10/09 OS fails.

> Note:
> (1) The problem does not occur if no zone clusters are installed on the cluster being upgraded.
> (2) The problem does not occur if there are zone clusters, but they exist only in the `Configured` state.

Workaround: Add two files to the Oracle Solaris installation miniroot image on the JumpStart server, as shown in the following steps. In these instructions, server is the name of the JumpStart server that you use to install the upgrade. Perform all steps as superuser.

Perform Step 1 from one node of the cluster that is to be upgraded.

1. Copy the zone-cluster `config.xml` and `platform.xml` files from a node that is to be upgraded to a location on the JumpStart server.

```
# cp /usr/lib/brand/cluster/config.xml.upgrade /net/server/some_dir/config.xml
# cp /usr/lib/brand/cluster/platform.xml /net/server/some_dir/platform.xml
```

Perform Steps 2-6 on the JumpStart server.

2. Create a new Oracle Solaris installation image.

```
# cd path_to_Solaris_install_image/Solaris_10/Tools
# ./setup_install_server path_to_new_Solaris_install_image
# cd
```

3. Unpack the miniroot.

```
# /boot/solaris/bin/root_archive unpackmedia path_to_new_Solaris_install_image
path_to_miniroot
```

4. Put the zone-cluster brand files in the miniroot.

```
# mkdir path_to_miniroot/usr/lib/brand/cluster
# cp /net/server/some_dir/config.xml path_to_miniroot/usr/lib/brand/cluster
# cp /net/server/some_dir/platform.xml path_to_miniroot/usr/lib/brand/cluster
```

5. Pack the miniroot and place it in the new installation image.

```
# /boot/solaris/bin/root_archive packmedia path_to_new_Solaris_install_image
path_to_miniroot
```

6. Run the `addclient` command from the new installation image location for each cluster node to be upgraded.

```
# cd path_to_new_Solaris_install_image/Solaris_10/Tools
# ./addclient your_configuration_details
```

Perform this workaround just once for each JumpStart server and for each Oracle Solaris update to which you plan to upgrade.

Top

Setting the `num_zoneclusters` Property Causes a Core Dump on a Cluster That Was Upgraded from Sun Cluster 3.1 to Sun Cluster 3.2 1/09 Software (6846142)

Problem Summary: The property `num_zoneclusters` that was introduced in Sun Cluster 3.2 1/09 software defines the maximum number of zone clusters that can be created on the cluster. When a cluster is upgraded to Oracle Solaris Cluster 3.3 software from a version earlier than the Sun Cluster 3.2 1/09 release, you must set this property by using the `cluster` command

before you configure any zone clusters. However, if the cluster is upgraded to Oracle Solaris Cluster 3.3 from Sun Cluster 3.1 software, the process of setting the `num_zoneclusters` property generates a core dump.

Workaround: Set the properties `max_nodes` and `max_privatenets` before you set the `num_zoneclusters` property. Use the following procedure:

1) Reboot all nodes into non-cluster mode.
2) Use the `/usr/cluster/bin/clsetup` or `/usr/cluster/bin/cluster` commands to set the values of the properties `max_nodes` and `max_privatenets`.
3) Reboot all nodes into cluster mode.
4) Run the `/usr/cluster/bin/cluster` command to set the value of `num_zoneclusters`.
Use the following command-line syntax, where `172.16.0.0` is an example private-network address:

```
# cluster set-netprops \
-p private_netaddr="172.16.0.0" \
-p max_nodes=number_of_nodes_expected_in_the_cluster \
-p max_privatenets=number_of_networks_expected_in_the_cluster
```

5) Start the `clsetup` utility and make selections similar to those shown:

```
# clsetup
*** Main Menu ***


1) Change Network Addressing and Ranges for the Cluster Transport
2) Show Network Addressing and Ranges for the Cluster Transport

?) Help with menu options
q) Quit

    Option: 1

>>> Change Network Addressing and Ranges for the Cluster Transport <<<


Network addressing for the cluster transport is currently configured
as follows:

Private Network ===
 private_netaddr: 172.16.0.0
   private_netmask: 255.255.240.0
   max_nodes: 64
   max_privatenets: 10
   num_zoneclusters: 12


Do you want to change this configuration (yes/no) [yes]? yes

The default network address for the cluster transport is 172.16.0.0.

Do you want to use the default (yes/no) [yes]? yes

The default netmask for the cluster transport is 255.255.240.0.

Do you want to use the default (yes/no) [yes]? no

The combination of private netmask and network address will dictate
both the maximum number of nodes and private networks that can be
supported by a cluster. Given your private network address, this
program will generate a range of recommended private netmasks based on
the maximum number of nodes and private networks that you anticipate
for this cluster.

In specifying the anticipated number of maximum nodes and private
networks for this cluster, it is important that you give serious
consideration to future growth potential. While both the private
netmask and network address can be changed later, the tools for making
such changes require that all nodes in the cluster be booted into
noncluster mode.

Maximum number of nodes anticipated for future growth [64]? 64

Maximum number of private networks anticipated for future growth [10]? 10

Specify a netmask of 255.255.248.0 to meet anticipated future
requirements of 64 cluster nodes and 10 private networks.

To accommodate more growth, specify a netmask of 255.255.240.0 to
support up to 64 cluster nodes and 20 private networks.

What netmask do you want to use [255.255.248.0]? 255.255.248.0
Is it okay to proceed with the update (yes/no) [yes]? yes
```

Top

Oracle Solaris Cluster 3.3 Localization

Product Localization

Localization for certain components of Oracle Solaris Cluster 3.3 software is as follows:

Component	Localization
Software Command Line	Japanese, Simplified Chinese
Software GUI	French, Japanese, Simplified Chinese, Spanish
Online Help	French, Japanese, Simplified Chinese, Spanish
Man Pages	Japanese

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`

Top

Oracle Solaris Cluster 3.3 Documentation

Documentation Set

This section lists the collections of the Oracle Solaris Cluster 3.3 user documentation set.

Oracle Solaris Cluster 3.3 Software Manuals

Part Number	Book Title
821-1261	Oracle Solaris Cluster 3.3 Documentation Center
821-1254	Oracle Solaris Cluster Concepts Guide
821-1258	Oracle Solaris Cluster Data Services Developer's Guide
821-1260	Oracle Solaris Cluster Data Services Planning and Administration Guide
821-1259	Oracle Solaris Cluster Error Messages Guide
821-1253	Oracle Solaris Cluster Overview
821-1573	Oracle Solaris Cluster Quick Reference
821-1255	Oracle Solaris Cluster Software Installation Guide
821-1257	Oracle Solaris Cluster System Administration Guide
821-1256	Oracle Solaris Cluster Upgrade Guide

Top

Oracle Solaris Cluster 3.3 Reference Manuals

Part Number	Book Title
821-1263	Oracle Solaris Cluster Reference Manual
821-1264	Oracle Solaris Cluster Data Services Reference Manual
821-1554	Oracle Solaris Cluster Quorum Server Reference Manual

Top

Oracle Solaris Cluster 3.3 Data Service Manuals (SPARC Platform Edition)

Part Number	Book Title
821-1572	Oracle Solaris Cluster Data Service for Agfa IMPAX Guide
821-1571	Oracle Solaris Cluster Data Service for Apache Guide
821-1522	Oracle Solaris Cluster Data Service for Apache Tomcat Guide
821-1523	Oracle Solaris Cluster Data Service for DHCP Guide
821-1524	Oracle Solaris Cluster Data Service for DNS Guide
821-1525	Oracle Solaris Cluster Data Service for Informix Guide
821-1526	Oracle Solaris Cluster Data Service for Kerberos Guide
821-1527	Oracle Solaris Cluster Data Service for MaxDB Guide
821-1528	Oracle Solaris Cluster Data Service for MySQL Guide
821-1932	Oracle Solaris Cluster Data Service for MySQL Cluster Guide
821-1530	Oracle Solaris Cluster Data Service for NFS Guide
821-1531	Oracle Solaris Cluster Data Service for Oracle Guide
821-1532	Oracle Solaris Cluster Data Service for Oracle Application Server Guide
821-2132	Oracle Solaris Cluster Data Service for Oracle Business Intelligence Enterprise Edition Guide
821-1533	Oracle Solaris Cluster Data Service for Oracle E-Business Suite Guide
821-1541	Oracle Solaris Cluster Data Service for Oracle Grid Engine Guide
821-1262	Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide
821-1687	Oracle Solaris Cluster Data Service for Oracle VM Server for SPARC Guide
821-2109	Oracle Solaris Cluster Data Service for PeopleSoft Enterprise Guide
821-1534	Oracle Solaris Cluster Data Service for PostgreSQL Guide
821-1535	Oracle Solaris Cluster Data Service for Samba Guide
821-1536	Oracle Solaris Cluster Data Service for SAP Guide
821-1537	Oracle Solaris Cluster Data Service for SAP liveCache Guide
821-1538	Oracle Solaris Cluster Data Service for SAP Web Application Server Guide
821-1539	Oracle Solaris Cluster Data Service for Siebel Guide
821-1540	Oracle Solaris Cluster Data Service for Solaris Zones Guide
821-1542	Oracle Solaris Cluster Data Service for Sun Java System Application Server EE (HADB) Guide
821-1543	Oracle Solaris Cluster Data Service for Sun Java System Application Server Guide
821-1544	Oracle Solaris Cluster Data Service for Sun Java System Message Queue Guide
821-1546	Oracle Solaris Cluster Data Service for Sun Java System Web Server Guide

821-1547	Oracle Solaris Cluster Data Service for SWIFTAlliance Access Guide
821-1548	Oracle Solaris Cluster Data Service for SWIFTAlliance Gateway Guide
821-1549	Oracle Solaris Cluster Data Service for Sybase ASE Guide
821-1550	Oracle Solaris Cluster Data Service for WebLogic Server Guide
821-1551	Oracle Solaris Cluster Data Service for WebSphere Message Broker Guide
821-1552	Oracle Solaris Cluster Data Service for WebSphere MQ Guide

Top

Oracle Solaris Cluster 3.3 Data Service Manuals (x86 Platform Edition)

Part Number	Book Title
821-1571	Oracle Solaris Cluster Data Service for Apache Guide
821-1522	Oracle Solaris Cluster Data Service for Apache Tomcat Guide
821-1523	Oracle Solaris Cluster Data Service for DHCP Guide
821-1524	Oracle Solaris Cluster Data Service for DNS Guide
821-1525	Oracle Solaris Cluster Data Service for Informix Guide
821-1526	Oracle Solaris Cluster Data Service for Kerberos Guide
821-1527	Oracle Solaris Cluster Data Service for MaxDB Guide
821-1528	Oracle Solaris Cluster Data Service for MySQL Guide
821-1932	Oracle Solaris Cluster Data Service for MySQL Cluster Guide
821-1530	Oracle Solaris Cluster Data Service for NFS Guide
821-1531	Oracle Solaris Cluster Data Service for Oracle Guide
821-1532	Oracle Solaris Cluster Data Service for Oracle Application Server Guide
821-1541	Oracle Solaris Cluster Data Service for Oracle Grid Engine Guide
821-1262	Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide
821-1534	Oracle Solaris Cluster Data Service for PostgreSQL Guide
821-1535	Oracle Solaris Cluster Data Service for Samba Guide
821-1536	Oracle Solaris Cluster Data Service for SAP Guide
821-1537	Oracle Solaris Cluster Data Service for SAP liveCache Guide
821-1538	Oracle Solaris Cluster Data Service for SAP Web Application Server Guide
821-1540	Oracle Solaris Cluster Data Service for Solaris Zones Guide
821-1542	Oracle Solaris Cluster Data Service for Sun Java System Application Server EE (HADB) Guide
821-1543	Oracle Solaris Cluster Data Service for Sun Java System Application Server Guide
821-1544	Oracle Solaris Cluster Data Service for Sun Java System Message Queue Guide
821-1546	Oracle Solaris Cluster Data Service for Sun Java System Web Server Guide
821-1549	Oracle Solaris Cluster Data Service for Sybase ASE Guide
821-1550	Oracle Solaris Cluster Data Service for WebLogic Server Guide
821-1551	Oracle Solaris Cluster Data Service for WebSphere Message Broker Guide
821-1552	Oracle Solaris Cluster Data Service for WebSphere MQ Guide

Top

Oracle Solaris Cluster 3.3 Hardware Collection (SPARC Platform Edition)

Part Number	Book Title
821-1555	Oracle Solaris Cluster 3.3 Hardware Administration Manual
821-1556	Oracle Solaris Cluster 3.3 With Network-Attached Storage Devices Manual
821-1557	Oracle Solaris Cluster 3.3 With SCSI JBOD Storage Device Manual
821-1558	Oracle Solaris Cluster 3.3 With StorageTek RAID Arrays Manual
821-1559	Oracle Solaris Cluster 3.3 With StorageTek 2540 RAID Arrays Manual
821-1560	Oracle Solaris Cluster 3.3 With Sun StorEdge 3310 or 3320 SCSI RAID Array Manual
821-1561	Oracle Solaris Cluster 3.3 With Sun StorEdge 3510 or 3511 FC RAID Array Manual
821-1562	Oracle Solaris Cluster 3.3 With Sun StorEdge 3900 Series or Sun StorEdge 6900 Series System Manual
821-1563	Oracle Solaris Cluster 3.3 With Sun StorEdge 6120 Array Manual
821-1564	Oracle Solaris Cluster 3.3 With Sun StorEdge 6130 Array Manual
821-1565	Oracle Solaris Cluster 3.3 With Sun StorEdge 6320 System Manual
821-1566	Oracle Solaris Cluster 3.3 With Sun StorEdge 9900 Series Storage Device Manual
821-1567	Oracle Solaris Cluster 3.3 With StorEdge A1000 Array, Netra st A1000 Array, or StorEdge A3500 System Manual
821-1568	Oracle Solaris Cluster 3.3 With Fibre Channel JBOD Storage Device Manual
821-1569	Oracle Solaris Cluster 3.3 With Sun StorEdge T3 or T3+ Array Manual
821-1570	Oracle Solaris Cluster 3.3 With Sun StorEdge 3500FC System Manual

Top

Oracle Solaris Cluster 3.3 Hardware Collection (x86 Platform Edition)

Part Number	Book Title
821-1555	Oracle Solaris Cluster 3.3 Hardware Administration Manual
821-1556	Oracle Solaris Cluster 3.3 With Network-Attached Storage Devices Manual
821-1557	Oracle Solaris Cluster 3.3 With SCSI JBOD Storage Device Manual
821-1558	Oracle Solaris Cluster 3.3 With StorageTek RAID Arrays Manual
821-1559	Oracle Solaris Cluster 3.3 With StorageTek 2540 RAID Arrays Manual
821-1561	Oracle Solaris Cluster 3.3 With Sun StorEdge 3510 or 3511 FC RAID Array Manual
821-1563	Oracle Solaris Cluster 3.3 With Sun StorEdge 6120 Array Manual
821-1564	Oracle Solaris Cluster 3.3 With Sun StorEdge 6130 Array Manual
821-1565	Oracle Solaris Cluster 3.3 With Sun StorEdge 6320 System Manual
821-1566	Oracle Solaris Cluster 3.3 With Sun StorEdge 9900 Series Storage Device Manual

Top

Oracle Solaris Cluster 3.3 Documentation Issues

Documentation Issues

This section discusses errors or omissions for documentation, online help, or man pages in the Oracle Solaris Cluster 3.3 release.

- System Administration Guide
- Data Service for Apache Tomcat Guide
- Data Service for Oracle Guide
- Data Service for Oracle Real Application Clusters Guide
- Network-Attached Storage Manual
- Man Pages
- Online Help

Top

Data Service for Oracle Real Application Clusters Guide

This section discusses errors and omissions in the Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide.

CLI Support to Configure Oracle RAC Instance Proxy Resource Group

Support is added to the Oracle Solaris Cluster maintenance commands to configure an Oracle Real Application Clusters (Oracle RAC) instance proxy resource group when Oracle ASM is not also configured. This support is delivered starting in the following HA-Oracle patch versions:

- SPARC: 145335-01
- x86: 145336-01

For supporting configuration instructions, see Configuring an Oracle RAC Instance Proxy Resource Group Without ASM (CLI) .

Procedures for removing the Grid Infrastructure resource are also provided. See Removing an Oracle Grid Infrastructure Resource.

Top

Configuring an Oracle RAC Instance Proxy Resource Group Without Oracle ASM (CLI)

For Oracle 11g release 2, a patch and additional configuration procedures are required to support using Oracle Solaris Cluster maintenance commands to configure an Oracle RAC instance proxy resource group without Oracle ASM. Perform the following tasks to add this support:

Task	Additional Information
1. Apply at least the following minimum HA-Oracle patch version on each node that will run Oracle 11g release 2: • SPARC: 145335-01 • x86: 145336-01	
2. Configure the disk set or disk group to use.	Follow one of the following procedures in the Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide: • If you are using Solaris Volume Manager for Sun Cluster, create a multi-owner disk set. Follow procedures in How to Create a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster for the Oracle RAC Database . • If you are using Veritas Volume Manager with the cluster feature, create a shared disk group. Follow procedures in How to Create a VxVM Shared-Disk Group for the Oracle RAC Database .

3. Perform procedures in Creating an Oracle Grid Infrastructure Resource.	When you configure Oracle RAC 11g release 2, this procedure replaces How to Create an Oracle Clusterware Resource for Interoperation With Sun Cluster in Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide. You can alternatively perform this task by using the Oracle Solaris Cluster Manager GUI or `clsetup` wizards, as described in How to Enable Oracle Solaris Cluster and Oracle 10g Release 2 or 11g Oracle Clusterware to Interoperate . If you do, and a patch for CR 6995923 is not yet available, you must first perform Step 1 through Step 3 of this procedure, to create the prerequisite Oracle Grid Infrastructure resource, before you run the wizard. The wizard does not currently create a missing Oracle Grid Infrastructure resource for you.
4. Perform the appropriate CLI procedure in Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide to create Oracle Solaris Cluster resources for interoperation with Oracle .	When an offline-restart dependency is set, the VALIDATE method of the `SUNW.scalable_rac_server_proxy` resource transparently modifies the appropriate Grid Infrastructure database resource as follows: • `ora.dbname.db` --> [add hard-start dependency] --> `sun.scal-dg1-rs` • `ora.dbname.db` --> [add hard-start dependency] --> `sun.scal-dg2-rs`

Top

Creating an Oracle Grid Infrastructure Resource

Perform this procedure to manually create an Oracle Grid Infrastructure resource that proxies the Oracle Solaris Cluster `SUNW.ScalDeviceGroup` or `SUNW.ScalMountPoint` resource. In this example procedure, the Oracle Grid Infrastructure resource is named `sun.`*`resource`*. You configure `sun.`*`resource`* to ensure that the corresponding Oracle database is not started until `sun.`*`resource`* is online. The `sun.`*`resource`* resource comes online only if the corresponding `SUNW.ScalDeviceGroup` or `SUNW.ScalMountPoint` resource is online. And the `SUNW.ScalDeviceGroup` or `SUNW.ScalMountPoint` resource only comes online if the actual volume-manager disk set or disk group or the mount point is online.

Top

How to Create an Oracle Grid Infrastructure Resource

Perform this procedure on one node of the cluster.

1. Become superuser.

2. Create the Oracle Grid Infrastructure **`sun.storage_proxy.type`** resource type.

In this example procedure, the `/export/home/oinstall/app/11.2.0/grid` directory is the Oracle Grid Infrastructure home directory.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl \
add type sun.storage_proxy.type \
-basetype cluster_resource \
-attr \
"ATTRIBUTE=ACTION_SCRIPT,TYPE=string", \
"ATTRIBUTE=HOSTING_MEMBERS,TYPE=string", \
"ATTRIBUTE=CARDINALITY,TYPE=string", \
"ATTRIBUTE=PLACEMENT,TYPE=string", \
"ATTRIBUTE=SCRIPT_TIMEOUT,TYPE=int", \
"ATTRIBUTE=RESTART_ATTEMPTS,TYPE=int", \
"ATTRIBUTE=ACL,TYPE=string", \
"ATTRIBUTE=VERSION,TYPE=string"
```

3. Create the Oracle Solaris Cluster **`sun.`*`resource`*** resource of type **`sun.storage_proxy.type`**.

The Oracle Grid Infrastructure resource name uses the form `sun.`*`sc-resource`*, where *`sc-resource`* is the name of the `SUNW.ScalDeviceGroup` or `SUNW.ScalMountPoint` resource. For example, the Oracle Solaris Cluster resource `scal-dg1-rs` is identified as `sun.scal-dg1-rs` in Oracle Grid Infrastructure.

You must also set the following attributes to values that are appropriate for your cluster:

- `HOSTING_MEMBERS` = The list of nodes in the cluster membership
- `CARDINALITY` = The number of nodes in the cluster membership

To create the Oracle Grid Infrastructure `storage_proxy` resource, execute the following command, substituting appropriate values for your cluster:

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl add resource sun.scal-dg1-rs \
-type sun.storage_proxy.type \
-attr "ACTION_SCRIPT='/opt/SUNWscor/dsconfig/bin/scproxy_crs_action' \
CARDINALITY='3' \
SCRIPT_TIMEOUT='20' \
PLACEMENT='restricted' \
RESTART_ATTEMPTS='60' \
HOSTING_MEMBERS='pnsx1 pnsx2 pnsx3' \
VERSION='1' "
```

4. Determine the DBA group of the Oracle Grid Infrastructure installation.

In this example procedure, the `/export/home/oinstall/app/11.2.0/grid` directory is the Oracle Grid Infrastructure home directory.

```
# echo `/export/home/oinstall/app/11.2.0/grid/bin/osdbagrp`
griddba
```

5. Set the primary group of the Oracle Grid Infrastructure **`storage_proxy`** resource to the group determined in Step 4.

In this example procedure, the `/export/home/oinstall/app/11.2.0/grid` directory is the Oracle Grid Infrastructure home directory.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl setperm resource sun.scal-dg1-rs -g "griddba"
```

6. Determine the DBA group of the Oracle Database Software installation.

In this example procedure, the `/export/home/oinstall/app/oracle/product/11.2.0/db` directory is the Oracle Database Software home directory.

```
# echo `/export/home/oinstall/app/oracle/product/11.2.0/db/bin/osdbagrp`
dba
```

7. Set the group permissions of the Oracle Grid Infrastructure **`storage_proxy`** resource to the group determined in Step 6.

Omit this step if the Oracle Grid Infrastructure installation DBA group determined in Step 4 and the Oracle Database Software installation DBA group determined in Step 6 are the same DBA group.

In this example procedure, the `/export/home/oinstall/app/11.2.0/grid` directory is the Oracle Grid Infrastructure home directory.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl setperm resource sun.scal-dg1-rs -u "group:dba:r-x"
```

8. Bring online the Oracle Grid Infrastructure **`storage_proxy`** resource.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl start resource sun.scal-dg1-rs
CRS-2672: Attempting to start 'sun.scal-dg1-rs' on 'pnsx2'
CRS-2672: Attempting to start 'sun.scal-dg1-rs' on 'pnsx3'
CRS-2672: Attempting to start 'sun.scal-dg1-rs' on 'pnsx1'
CRS-2676: Start of 'sun.scal-dg1-rs' on 'pnsx2' succeeded
CRS-2676: Start of 'sun.scal-dg1-rs' on 'pnsx3' succeeded
CRS-2676: Start of 'sun.scal-dg1-rs' on 'pnsx1' succeeded
```

Top

Removing an Oracle Grid Infrastructure Resource

This section contains the following procedures to remove an Oracle Grid Infrastructure resource.

- How to Remove a Dependency
- How to Delete the `sun.resource` Resource

Top

How to Remove a Dependency

This procedure shows how to set the offline-restart dependency to remove a dependency.

1. Become superuser.

2. Display the current start dependency that the database has on the Oracle Grid Infrastructure **`storage_proxy`** resource.

In this example procedure, the `/export/home/oinstall/app/11.2.0/grid` directory is the Oracle Grid Infrastructure home directory, the database name is `testdb`, and the Oracle Grid Infrastructure `storage_proxy` resource name is `sun.scal-dg1-rs`.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl stat res ora.testdb.db -p | grep START_DEPENDENCIES
START_DEPENDENCIES=hard(sun.scal-dg1-rs)
weak(type:ora.listener.type,global:type:ora.scan_listener.type,uniform:ora.ons,uniform:ora.eons)

# clresource show -p resource_dependencies_offline_restart rac-server-proxy-rs
=== Resources ===

Resource: rac-server-proxy-rs
Resource_dependencies_offline_restart: crs-fw-rs scal-dg1-rs

— Standard and extension properties —
```

3. Remove the offline-restart dependency on the **`SUNW.ScalDeviceGroup`** or **`SUNW.ScalMountPoint`** resource from the Oracle RAC instance proxy resource.

This command clears the dependencies that the Oracle Grid Infrastructure database resource has on the Oracle Grid Infrastructure `storage_proxy` resource. Note that the command includes the minus (-) symbol.

```
# clresource set -p resource_dependencies_offline_restart-=scal-dg1-rs rac-server-proxy-rs
```

4. Verify that the start dependency on the Oracle Grid Infrastructure resource is removed.

In this example procedure, the `/export/home/oinstall/app/11.2.0/grid` directory is the Oracle Grid Infrastructure home directory.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl stat res ora.testdb.db -p | grep START_DEPENDENCIES
START_DEPENDENCIES=weak(type:ora.listener.type,global:type:ora.scan_listener.type,uniform:ora.ons,uniform:ora.eons

# clresource show -p resource_dependencies_offline_restart rac-server-proxy-rs

=== Resources ===

Resource: rac-server-proxy-rs
Resource_dependencies_offline_restart: crs-fw-rs

— Standard and extension properties —
```

Top

How to Delete the **`sun.resource`** Resource

1. Become superuser.

2. Ensure that the dependency is removed, as described in How to Remove a Dependency, and that **sun.resource** is stopped.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl stop res sun.scal-dg1-rs
CRS-2673: Attempting to stop 'sun.scal-dg1-rs' on 'pnsx3'
CRS-2673: Attempting to stop 'sun.scal-dg1-rs' on 'pnsx1'
CRS-2673: Attempting to stop 'sun.scal-dg1-rss' on 'pnsx2'
CRS-2677: Stop of 'sun.scal-dg1-rs' on 'pnsx3' succeeded
CRS-2677: Stop of 'sun.scal-dg1-rs' on 'pnsx1' succeeded
CRS-2677: Stop of 'sun.scal-dg1-rs' on 'pnsx2' succeeded
```

3. Delete **sun.*resource***.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl delete res sun.scal-dg1-rs
```

4. Verify that **sun.*resource*** is deleted.

```
# /export/home/oinstall/app/11.2.0/grid/bin/crsctl stat res sun.scal-dg1-rs -p
CRS-210: Could not find resource 'sun.scal-dg1-rs'.
```

Top

System Administration Guide

Extraneous Information in "Adding a Quorum Device"

The section "Adding a Quorum Device" refers to the use of Oracle's Sun Storage 7000 Unified Storage System as a NAS quorum device. Support for Oracle's Sun Storage 7000 Unified Storage System as a NAS device has not been enabled in the 3.3 release. Ignore the information.

Top

Data Service for Apache Tomcat Guide

Changes With Apache Tomcat 5.5.28 and 6.0.29

The following information additions and changes support the `wget` probe algorithm for Oracle Solaris Cluster HA for Apache Tomcat (HA for Apache Tomcat) using Apache Tomcat versions starting with 5.5.28 and 6.0.29. This functionality is delivered in the following Oracle Solaris Cluster 3.3 patches for HA for Apache Tomcat:

- SPARC: 145341-01
- x86: 145342-01

- New Restriction For Apache Tomcat 5.5.28 and 6.0.29
- New Requirement for Apache Tomcat 5.5.28 and 6.0.29
- Installing and Configuring Apache Tomcat
- How to Register and Configure Solaris Cluster HA for Apache Tomcat as a Failover Data Service
- How to Register and Configure Solaris Cluster HA for Apache Tomcat as a Multiple Masters Data Service
- How to Register and Configure Solaris Cluster HA for Apache Tomcat as a Scalable Data Service
- Structure of the Apache Tomcat HA parameter file
- Strategy to Chose the TestCmd and the ReturnString Variable
- Probing Algorithm and Functionality
- Deployment Example: Installing Apache Tomcat in the Global Zone
- Example: Modifying the Apache Tomcat Configuration
- Deployment Example: Installing Apache Tomcat in a Failover Zone
- Example: Installing the Apache Tomcat Software on Shared Storage
- Deployment Example: Installing Apache Tomcat in a Non-Global Zone
- Example: Modifying the Apache Tomcat Configuration File

New Restriction For Apache Tomcat 5.5.28 and 6.0.29

Starting with Apache Tomcat versions 5.5.28 and 6.0.29, you must use the `wget` probe algorithm. This bypasses the normal probe, enabling the TestUrl parameter to work correctly. You can optionally use the `wget` probe algorithm for earlier versions as well.

Top

New Requirement for Apache Tomcat 5.5.28 and 6.0.29

To support the `wget` probe algorithm, the `wget` command must be available from every node of the cluster.

Top

Installing and Configuring Apache Tomcat

In the bullet list of considerations to determine how Apache Tomcat will be deployed with Oracle Solaris Cluster, you must also determine which Apache Tomcat version to deploy.

Top

How to Register and Configure Solaris Cluster HA for Apache Tomcat as a Failover Data Service

The following replaces Step 1:

1. Prepare the parameter file, which is required by HA for Apache Tomcat.

```
# cd /opt/SUNWsctomcat/bin
# cp pfile desired-place
```

Edit the parameter file `pfile` and follow the comments within that file. For example:

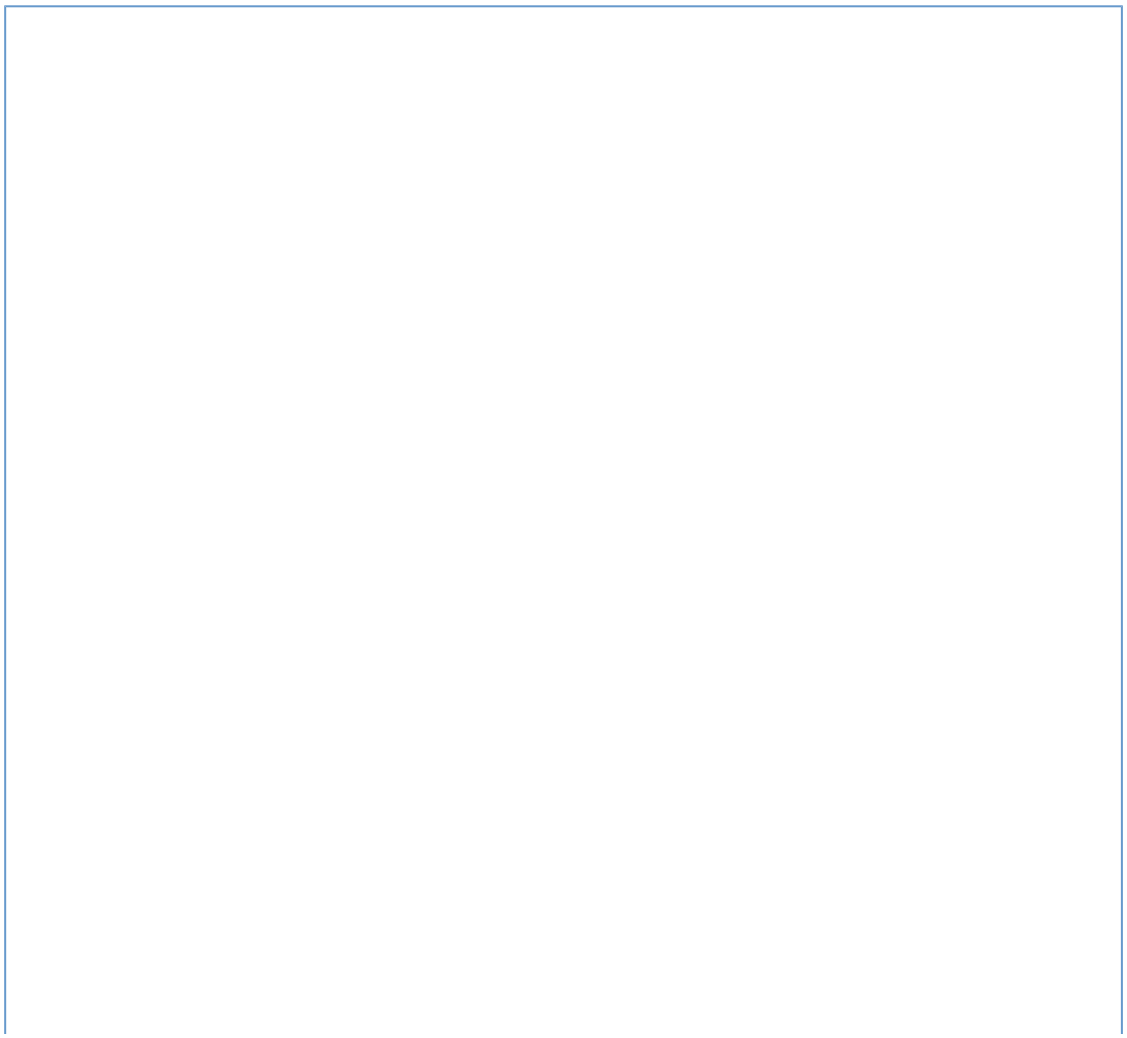

```
#!/usr/bin/ksh
#
# CDDL HEADER START
#
# The contents of this file are subject to the terms of the
# Common Development and Distribution License (the License).
# You may not use this file except in compliance with the License.
#
# You can obtain a copy of the license at usr/src/CDDL.txt
# or http://www.opensolaris.org/os/licensing.
# See the License for the specific language governing permissions
# and limitations under the License.
#
# When distributing Covered Code, include this CDDL HEADER in each
# file and include the License file at usr/src/CDDL.txt.
# If applicable, add the following below this CDDL HEADER, with the
# fields enclosed by brackets \[\] replaced with your own identifying
# information: Portions Copyright \[yyyy\] \[name of copyright owner\]
#
# CDDL HEADER END
#
#
# Copyright (c) 2006, 2010, Oracle and/or its affiliates. All rights reserved.
#

#ident  "@(#)pfile.ksh  1.6     10/08/16"

# Set the Apache Tomcat specific environment variables which the start, stop
# and check fuctions will use
#
# EnvScript     Script to set runtime environment for tomcat
# User          Apache Tomcat User
# Basepath      Absolute path to Tomcat Home directory i.e. 4.x CATALINA_HOME
#               or TOMCAT_HOME for 3.x
# Host          Hostname to test Apache Tomcat
# Port          Port where Apache Tomcat is configured
# TestCmd       Apache Tomcat test command
#
# Some versions of Apache Tomcat do not react on the probe algorithm with the
# http get command. Because of this we provide an alternative probe method
# which uses wget. If wget is not installed on your system you must install it
# manually and provide the WgetPath variable. If the WgetPath variable is not set,
# it defaults to /usr/sfw/bin/wget.
# You can use the wget algorithm for https probe as well which are not possible
# with the http get command.
#
# Use the variables TestUrl and WgetPath as a replacement for:
# Port, Host, TestCmd
# You can specify only one set of variables.
#
# TestUrl       Url where Tomcat serve a web site. This can be done either
#               by the http or https  protocol
#               Format: http://hostname:port/startpage
#               Example: http://localhost:8080/
#               You might want to include options here like:
#               "--no-cookies http://localhost:8080/"
#               The requirement here is that wget -O - ${TestUrl} produces
#               the output containing the ReturnString.
# WgetPath      Abolute pat to wget, leaving this variable empty defaults to:
#               /usr/sfw/bin/wget
# Startwait     Sleeping $Startwait seconds after completion of the
#               start command
# ReturnString  This string must be present in the output of the http get command
#               or in the output of the wget command.
EnvScript=
User=
Basepath=
Host=
Port=8080
TestCmd="get /index.jsp"
TestUrl=
WgetPath=
ReturnString="CATALINA"
Startwait=20
```

The following is an example for Apache Tomcat 4.1.24:

```
EnvScript=/global/mnt1/tomtest/env.ksh
User=tomtest
Basepath=/global/mnt1/tomtest/jakarta-tomcat-4.1.24
Host=tomcat-lh
Port=8080
TestCmd="get /index.jsp"
TestUrl=
WgetPath=
ReturnString="CATALINA"
Startwait=20
```

This example for Apache Tomcat 4.1.24 configures a standalone web server with the default start page `index.jsp`. The start page contains the string `CATALINA`.

To use the alternative `wget` probe, assuming `wget` is available under `/usr/sfw/bin/wget`, instead set the following variables:

```
EnvScript=/global/mnt1/tomtest/env.ksh
User=tomtest
Basepath=/global/mnt1/tomtest/jakarta-tomcat-4.1.24
Host=
Port=
TestCmd=
TestUrl=”http://tomcat-lh:8080/”
WgetPath=
ReturnString="CATALINA"
Startwait=20
```

Note
The parameter files must be available on every node that can host the Apache Tomcat data service. For a scalable or multiple-masters configuration, store the parameter files on the local file system of every node or non-global zone or on the shared storage. The parameter files must not differ for any instance of Apache Tomcat on the various nodes.

Repeat Step 1 for every Apache Tomcat instance you need.

Top

How to Register and Configure Solaris Cluster HA for Apache Tomcat as a Multiple Masters Data Service

The following replaces Step 2:

2. Prepare the parameter file, which is required by HA for Apache Tomcat.

```
# cd /opt/SUNWsctomcat/util
# cp sctomcat_config desired-place
```

Edit the parameter file `pfile` and follow the comments within that file. For example:

```
#!/usr/bin/ksh
#
# CDDL HEADER START
#
# The contents of this file are subject to the terms of the
# Common Development and Distribution License (the License).
# You may not use this file except in compliance with the License.
#
# You can obtain a copy of the license at usr/src/CDDL.txt
# or http://www.opensolaris.org/os/licensing.
# See the License for the specific language governing permissions
# and limitations under the License.
#
```

```
# When distributing Covered Code, include this CDDL HEADER in each
# file and include the License file at usr/src/CDDL.txt.
# If applicable, add the following below this CDDL HEADER, with the
# fields enclosed by brackets \[\] replaced with your own identifying
# information: Portions Copyright \[yyyy\] \[name of copyright owner\]
#
# CDDL HEADER END
#

#
# Copyright (c) 2006, 2010, Oracle and/or its affiliates. All rights reserved.
#

#ident  "@(#)pfile.ksh  1.6     10/08/16"

# Set the Apache Tomcat specific environment variables which the start, stop
# and check fuctions will use
#
# EnvScript     Script to set runtime environment for tomcat
# User          Apache Tomcat User
# Basepath      Absolute path to Tomcat Home directory i.e. 4.x CATALINA_HOME
#               or TOMCAT_HOME for 3.x
# Host          Hostname to test Apache Tomcat
# Port          Port where Apache Tomcat is configured
# TestCmd       Apache Tomcat test command
#
# Some versions of Apache Tomcat do not react on the probe algorithm with the
# http get command. Because of this we provide an alternative probe method
# which uses wget. If wget is not installed on your system you must install it
# manually and provide the WgetPath variable. If the WgetPath variable is not set,
# it defaults to /usr/sfw/bin/wget.
# You can use the wget algorithm for https probe as well which are not possible
# with the http get command.
#
# Use the variables TestUrl and WgetPath as a replacement for:
# Port, Host, TestCmd
# You can specify only one set of variables.
#
# TestUrl       Url where Tomcat serve a web site. This can be done either
#               by the http or https  protocol
#               Format: http://hostname:port/startpage
#               Example: http://localhost:8080/
#               You might want to include options here like:
#               "--no-cookies http://localhost:8080/"
#               The requirement here is that wget -O - ${TestUrl} produces
#               the output containing the ReturnString.
# WgetPath      Abolute pat to wget, leaving this variable empty defaults to:
#               /usr/sfw/bin/wget
# Startwait     Sleeping $Startwait seconds after completion of the
#               start command
# ReturnString  This string must be present in the output of the http get command
#               or in the output of the wget command.
EnvScript=
User=
Basepath=
Host=
Port=8080
TestCmd="get /index.jsp"
TestUrl=
WgetPath=
```

```
ReturnString="CATALINA"
Startwait=20
```

The following is an example for Apache Tomcat 4.1.24.

```
EnvScript=/tomcat/env.ksh
User=tomcat
Basepath=/tomcat/jakarta-tomca-4.1.24
Host=localhost
Port=8080
TestCmd="get /index.jsp"
TestUrl=
WgetPath=
ReturnString="CATALINA"
Startwait=20
```

This example for Apache Tomcat 4.1.24 configures a standalone web server with the default start page `index.jsp`. The start page contains the string `CATALINA`.

Depending on the selected probe method, the Host parameter or the IP alias in the TestUrl parameter depends on the location of the `pfile`:

- If the `pfile` is stored on the shared storage, you must use `localhost`.
- If the `pfile` is stored on the local storage, it can be either `localhost` or the node's hostname.

To use the alternative `wget` probe, assuming `wget` is available under `/usr/sfw/bin/wget`, set the following variables:

```
EnvScript=/global/mnt1/tomtest/env.ksh
User=tomtest
Basepath=/global/mnt1/tomtest/jakarta-tomcat-4.1.24
Host=
Port=
TestCmd=
TestUrl=”http://localhost:8080/”
WgetPath=
ReturnString="CATALINA"
Startwait=20
```

Note
The parameter files must be available on every node that can host the Apache Tomcat data service. For a scalable or multiple-masters configuration, store the parameter files on the local file system of every node or non-global zone or on the shared storage. The parameter files must not differ for any instance of Apache Tomcat on the various nodes.

Repeat this Step for every Apache Tomcat instance and every node you need.

Top

How to Register and Configure Solaris Cluster HA for Apache Tomcat as a Scalable Data Service

The following replaces Step 2:

2. Prepare the parameter file, which is required by Oracle Solaris Cluster HA for Apache Tomcat.

```
# cd /opt/SUNWsctomcat/bin
# cp pfile desired-place
```

Edit the parameter file `pfile` and follow the comments within that file. For example:

```
#!/usr/bin/ksh
#
# CDDL HEADER START
```

```
#
# The contents of this file are subject to the terms of the
# Common Development and Distribution License (the License).
# You may not use this file except in compliance with the License.
#
# You can obtain a copy of the license at usr/src/CDDL.txt
# or http://www.opensolaris.org/os/licensing.
# See the License for the specific language governing permissions
# and limitations under the License.
#
# When distributing Covered Code, include this CDDL HEADER in each
# file and include the License file at usr/src/CDDL.txt.
# If applicable, add the following below this CDDL HEADER, with the
# fields enclosed by brackets \[\] replaced with your own identifying
# information: Portions Copyright \[yyyy\] \[name of copyright owner\]
#
# CDDL HEADER END
#

#
# Copyright (c) 2006, 2010, Oracle and/or its affiliates. All rights reserved.
#

#ident  "@(#)pfile.ksh  1.6     10/08/16"

# Set the Apache Tomcat specific environment variables which the start, stop
# and check fuctions will use
#
# EnvScript     Script to set runtime environment for tomcat
# User          Apache Tomcat User
# Basepath      Absolute path to Tomcat Home directory i.e. 4.x CATALINA_HOME
#               or TOMCAT_HOME for 3.x
# Host          Hostname to test Apache Tomcat
# Port          Port where Apache Tomcat is configured
# TestCmd       Apache Tomcat test command
#
# Some versions of Apache Tomcat do not react on the probe algorithm with the
# http get command. Because of this we provide an alternative probe method
# which uses wget. If wget is not installed on your system you must install it
# manually and provide the WgetPath variable. If the WgetPath variable is not set,
# it defaults to /usr/sfw/bin/wget.
# You can use the wget algorithm for https probe as well which are not possible
# with the http get command.
#
# Use the variables TestUrl and WgetPath as a replacement for:
# Port, Host, TestCmd
# You can specify only one set of variables.
#
# TestUrl       Url where Tomcat serve a web site. This can be done either
#               by the http or https  protocol
#               Format: http://hostname:port/startpage
#               Example: http://localhost:8080/
#               You might want to include options here like:
#               "--no-cookies http://localhost:8080/"
#               The requirement here is that wget -O - ${TestUrl} produces
#               the output containing the ReturnString.
# WgetPath      Abolute pat to wget, leaving this variable empty defaults to:
#               /usr/sfw/bin/wget
# Startwait     Sleeping $Startwait seconds after completion of the
#               start command
# ReturnString  This string must be present in the output of the http get command
#               or in the output of the wget command.
EnvScript=
User=
Basepath=
Host=
Port=8080
TestCmd="get /index.jsp"
TestUrl=
WgetPath=
```

```
ReturnString="CATALINA"
Startwait=20
```

The following is an example for Apache Tomcat 4.1.24.

```
EnvScript=/tomcat/env.ksh
User=tomcat
Basepath=/tomcat/jakarta-tomca-4.1.24
Host=localhost
Port=8080
TestCmd="get /index.jsp"
TestUrl=
WgetPath=
ReturnString="CATALINA"
Startwait=20
```

To use the alternative `wget` probe, assuming `wget` is available under `/usr/sfw/bin/wget`, set the following variables:

```
EnvScript=/global/mnt1/tomtest/env.ksh
User=tomtest
Basepath=/global/mnt1/tomtest/jakarta-tomcat-4.1.24
Host=
Port=
TestCmd=
TestUrl=”http://localhost:8080/”
WgetPath=
ReturnString="CATALINA"
Startwait=20
```

This example for Apache Tomcat 4.1.24 configures a standalone web server with the default start page `index.jsp`. The start page contains the string `CATALINA`.

Depending on the selected probe method, the Host parameter or the IP alias in the TestUrl parameter depends on the location of the `pfile`.

- If the `pfile` is stored on the shared storage, you must use `localhost`.
- If the `pfile` is stored on the local storage, it can be either `localhost` or the node's hostname.

Note
Do not use the SharedAddress here.

Note
The parameter files must be available on every node that can host the Apache Tomcat data service. For a scalable or multiple-masters configuration, store the parameter files on the local file system of every node or on the shared storage. The parameter files must not differ for an instance of Apache Tomcat on the various nodes.

Repeat this Step for every Apache Tomcat instance and every node you need.

Top

Structure of the Apache Tomcat HA Parameter File

The following are additions to Table 6: Structure of the Apache Tomcat HA parameter file.

Variable	Explanation
TestUrl	This variable presents a complete `http` or `https` URL, which is passed to the `wget` command to test the Tomcat server's sanity.
WgetPath	This variable presents the absolute path to the `wget` binary, if the `wget` binary is not located under `/usr/sfw/bin/wget`.

In addition, TestUrl and WgetPath are added to the list of parameters that are used for starting and stopping Apache Tomcat.

Top

Strategy to Choose the TestCmd and the ReturnString Variable

The following replaces the section "Strategy to Chose the TestCmd and the ReturnString Variable".

Strategy to Choose the TestCmd or TestUrl Variable and the ReturnString Variable

The following alternatives are available to choose the ReturnString variable:

- Take the start page of your application and set the TestCmd to get */start_page*. If you picked the `wget` probe method, set TestUrl to `http://`*ipalias*`:`*port*`/start_page`. With this strategy, you are monitoring that Apache Tomcat is serving your application.
- Take the Apache Tomcat default start page and set the TestCmd to get */default-startpage*. If you picked the `wget` method, set TestUrl to `http://`*ipalias*`:`*port*`/`*default-startpage*. In this case, set the ReturnString to a string contained in the startpage. This string depends on the deployed Apache Tomcat version; for 3.x it is Tomcat home and for 4.x and 5.x it is CATALINA. With this strategy, you are monitoring that Apache Tomcat is serving its default application.
- Deploy a test application (which is not provided with the HA for Apache Tomcat agent) to Apache Tomcat. Set the TestCmd to get *startpage_of_the_application*. If you use the `wget` method, set TestUrl to `http://`*ipalias*`:`*port*`/`*startpage_of_the_application*. In this case, set the ReturnString to a string contained in the start page. With this strategy, you are monitoring that Apache Tomcat is serving your test application.
- If none of above is appropriate, set the TestCmd to get */a-page-that-does-not-exist*. In this case, set the ReturnString to a string contained in the Error Page. With this strategy, you are monitoring that Apache Tomcat is operational, because it registers that it must deliver a page that does not exist.

You can evaluate the different pages by connecting from a browser to *hostname*`:`*port* and specifying the different pages.

Top

Probing Algorithm and Functionality

In the steps that describe how the sanity of Apache Tomcat is monitored, the following replaces Step 2:

2. Pings the Host, which is configured in the Apache Tomcat HA parameter file, unless the TestUrl parameter is specified in the Tomcat parameter file.

Top

Deployment Example: Installing Apache Tomcat in the Global Zone

In the section "Software Configuration", the following is an additional expectation of the example configuration:

- The `wget` binary is available under `/usr/sfw/bin/wget`.

Top

Example: Modifying the Apache Tomcat Configuration

The following is an addition to Step 4:

As an alternative, you can use the following set of variables:

```
EnvScript=/global/mnt3/tomcat/env.ksh
User=tomcat
Basepath=/global/mnt3/tomcat/apache-tomcat-5.5.17
TestUrl="http://ha-host-1:8080/"
ReturnString="CATALINA"
Startwait=20
```

Top

Deployment Example: Installing Apache Tomcat in a Failover Zone

In the section "Software Configuration", the following is an additional expectation of the example configuration:

- The `wget` binary is available under `/usr/sfw/bin/wget`.

Top

Example: Installing the Apache Tomcat Software on Shared Storage

The following is an addition to Step 6:

As an alternative, you can use the following set of variables:

```
EnvScript=/global/mnt3/tomcat/env.ksh
User=tomcat
Basepath=/global/mnt3/tomcat/apache-tomcat-5.5.17
TestUrl="http://ha-host-1:8080/"
ReturnString="CATALINA"
Startwait=20
```

Top

Deployment Example: Installing Apache Tomcat in a Non-Global Zone

In the section "Software Configuration", the following is an additional expectation of the example configuration:

- The `wget` binary is available under `/usr/sfw/bin/wget`.

Top

Example: Modifying the Apache Tomcat Configuration File

The following is an addition to Step 3:

As an alternative, you can use the following set of variables:

```
EnvScript=/global/mnt3/tomcat/env.ksh
User=tomcat
Basepath=/global/mnt3/tomcat/apache-tomcat-5.5.17
TestUrl="http://ha-host-1:8080/"
ReturnString="CATALINA"
Startwait=20
```

Top

Data Service for Oracle Guide

This section discusses errors and omissions in the Oracle Solaris Cluster Data Service for Oracle Guide.

Support for Oracle Grid Infrastructure for Clusters SCAN Listener with 11g Release 2

When the HA for Oracle data service is deployed using Oracle 11g release 2 Grid Infrastructure for Clusters, the preferred listener for a single-instance database is the Grid Infrastructure for Clusters Single Client Access Name (SCAN) Listener. To leverage SCAN in the HA-Oracle configuration, you must modify the Oracle 11g release 2 single-instance database remote_listener parameter to reflect the SCAN name and port number.

The following example shows the commands and sample output to configure the SCAN name and port.

```
# srvctl config scan
SCAN name: freak-scan-lh, Network: 1/10.11.188.0/255.255.255.0/bge0
SCAN VIP name: scan1, IP: /vzfreak1a/10.11.188.60

# srvctl config scan_listener
SCAN Listener LISTENER_SCAN1 exists. Port: TCP:1521
#

SQL> show parameters listener

NAME          TYPE   VALUE
------------------------------------ ----------- ------------------------------
listener_networks       string
local_listener        string
(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=freak-1)(PORT=1521))))
remote_listener        string


SQL>
SQL> alter system set remote_listener = 'freak-scan-lh:1521' scope=both;

System altered.
SQL> show parameters listener

NAME          TYPE   VALUE
------------------------------------ ----------- ------------------------------
listener_networks       string
local_listener        string
(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=freak-1)(PORT=1521))))
remote_listener        string      freak-scan-lh:1521
SQL>
```

If you do not want to use the SCAN listener, use the `SUNW.oracle_listener` resource type instead and configure it as usual.

Top

Support for Third-Party Volume Managers with Clustered Oracle ASM Disk Groups

For Oracle 11g release 2, a patch and additional configuration procedures are required to support third-party volume managers with clustered Oracle ASM disk groups. Perform the following tasks to add this support:

1. Apply the following minimum HA-Oracle patch on each node that will run Oracle 11g release 2 with clustered Oracle ASM:
   - SPARC: 145335-01
   - x86: 145336-01
2. Configure the disk set or disk group for Oracle ASM to use.
   Follow the appropriate procedure in the Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide:
   - If you are using Solaris Volume Manager for Sun Cluster, create a multi-owner disk set. Follow procedures in "How to Create a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster for the Oracle RAC Database" .
   - If you are using Veritas Volume Manager with the cluster feature, create a shared disk group. Follow procedures in "How to Create a VxVM Shared-Disk Group for the Oracle RAC Database" .
3. Perform procedures in "Creating an Oracle Grid Infrastructure Resource" .

Procedures for removing the Grid Infrastructure resource are also provided. See "Removing an Oracle Grid Infrastructure Resource" .

Top

Creating an Oracle Grid Infrastructure Resource

Perform this procedure to manually create a Grid Infrastructure resource that proxies the Oracle Solaris Cluster `SUNW.ScalDeviceGroup` resource. In this example procedure, the Grid Infrastructure resource is named `sun.resource`. You configure `sun.resource` to ensure that the corresponding Oracle ASM disk group is not mounted until `sun.resource` is online. The `sun.resource` resource comes online only if the corresponding `SUNW.ScalDeviceGroup` resource is online. And the `SUNW.ScalDeviceGroup` resource only comes online if the actual volume-manager disk set or disk group is online.

To ensure that the Oracle ASM disk group benefits from this dependency chain, after you define `sun.resource`, you modify the

appropriate Grid Infrastructure Oracle ASM disk-group resource so that the hard-start dependency includes `sun.resource`. Modifying the hard-start dependency of the Grid Infrastructure Oracle ASM disk-group resource can only be performed by the `SUNW.scalable_asm_diskgroup_proxy` resource by using the VALIDATE method. Therefore, you must set an offline-restart dependency between the `SUNW.scalable_asm_diskgroup_proxy` and `SUNW.ScalDeviceGroup` resources.

Top

How to Create an Oracle Grid Infrastructure Resource

This example procedure sets the following dependencies:

- A Grid Infrastructure resource, `sun.resource`, where `resource` corresponds to an Oracle Solaris Cluster `SUNW.ScalDeviceGroup` resource name:
  - `sun.scal-asmdg1-rs`
  - `sun.scal-asmdg2-rs`
- Oracle Solaris Cluster offline-restart dependencies:
  - `asm-data1-rs` --> [add offline-restart dependency] -> `scal-asmdg1-rs`
  - `asm-data2-rs` --> [add offline-restart dependency] -> `scal-asmdg2-rs`

When an offline-restart dependency is set, the VALIDATE method of the SUNW.scalable_asm_diskgroup_proxy resource will transparently modify the appropriate Grid Infrastructure ASM disk group resource as follows:

- Grid Infrastructure hard-start dependencies:
  - `ora.data1.dg` -> [add hard-start dependency] -> `sun.scal-asmdg1-rs`
  - `ora.data1.dg` -> [add hard-start dependency] -> `sun.scal-asmdg1-rs`

Perform this procedure on one node of the cluster.

1. Become superuser.
2. Create the **`sun.storage_proxy.type`** resource type.
   In this example procedure, the `/u01/app/11.2.0/grid/` directory is the Grid Infrastructure Oracle Home.

   ```
   bash-3.00# /u01/app/11.2.0/grid/bin/crsctl add type sun.storage_proxy.type -basetype cluster_resource
   ```

3. Create **`sun.resource`** of type **`sun.storage_proxy.type`**.

   > **Note**
   > Ensure that all attribute values are enclosed in single quotes. Otherwise, the VALIDATE method of the `SUNW.scalable_asm_diskgroup_proxy` resource will fail the validation.

   ```
   bash-3.00# /u01/app/11.2.0/grid/bin/crsctl add res sun.scal-asmdg1-rs -type sun.storage_proxy.type \
   -attr "ACTION_SCRIPT='/opt/SUNWscor/dsconfig/bin/scproxy_crs_action' \
   HOSTING_MEMBERS='pfreak1 pfreak2' \
   CARDINALITY='2' \
   PLACEMENT='restricted' \
   ACL='owner:root:rwx,pgrp:oinstall:rwx,other::r--' \
   SCRIPT_TIMEOUT='20' \
   RESTART_ATTEMPTS='60'"
   ```

`sun.scal-asmdg1-rs`	The `SUNW.ScalDeviceGroup` resource name.
`-type sun.storage_proxy.type`	Specify the `sun.storage_proxy.type` resource type.
`ACTION_SCRIPT`	Specifies the `/opt/SUNWscor/dsconfig/bin/scproxy_crs_action` script.
`HOSTING_MEMBERS`	Specifies the Oracle Solaris Cluster resource group's node list entries that contain the `SUNW.ScalDeviceGroup` resource.
`CARDINALITY`	Sets the number of hosts that are defined in HOSTING_MEMBERS.
`PLACEMENT`	Set to `restricted`.

ACL	Sets the owner equal to `root` and the group equal to the ACL group entry for the Oracle ASM disk group. The following command displays the ACL group entry:
	`bash-3.00# /u01/app/11.2.0/grid/bin/crsctl stat res ora.DATA1.dg -p | grep ACL=`
	`ACL=owner:oragrid:rwx,pgrp:oinstall:rwx,other::r--`
	The example output shows that `oinstall` is the group entry.

`SCRIPT_TIMEOUT`	Set to `20`.
`RESTART_ATTEMPTS`	Set to `60`.

4. Verify that **sun.*resource*** is correctly defined.
   Output is similar to the following:

```
bash-3.00# /u01/app/11.2.0/grid/bin/crsctl stat res sun.scal-asmdg1-rs -p
NAME=sun.scal-asmdg1-rs
TYPE=sun.storage_proxy.type
ACL=owner:root:rwx,pgrp:oinstall:rwx,other::r--
ACTION_FAILURE_TEMPLATE=
ACTION_SCRIPT=/opt/SUNWscor/dsconfig/bin/scproxy_crs_action
ACTIVE_PLACEMENT=0
AGENT_FILENAME=%CRS_HOME%/bin/scriptagent
AUTO_START=restore
CARDINALITY=2
CHECK_INTERVAL=60
DEFAULT_TEMPLATE=
DEGREE=1
DESCRIPTION=
ENABLED=1
FAILOVER_DELAY=0
FAILURE_INTERVAL=0
FAILURE_THRESHOLD=0
HOSTING_MEMBERS=pfreak1 pfreak2
LOAD=1
LOGGING_LEVEL=1
NOT_RESTARTING_TEMPLATE=
OFFLINE_CHECK_INTERVAL=0
PLACEMENT=restricted
PROFILE_CHANGE_TEMPLATE=
RESTART_ATTEMPTS=60
SCRIPT_TIMEOUT=20
SERVER_POOLS=
START_DEPENDENCIES=
START_TIMEOUT=0
STATE_CHANGE_TEMPLATE=
STOP_DEPENDENCIES=
STOP_TIMEOUT=0
UPTIME_THRESHOLD=1h
bash-3.00#
```

5. Set the offline-restart dependency to add a dependency.
   a. Display the current dependency.

```
bash-3.00# /u01/app/11.2.0/grid/bin/crsctl stat res ora.DATA1.dg -p | grep
START_DEPENDENCIES
START_DEPENDENCIES=hard(ora.asm) pullup(ora.asm)
bash-3.00# clresource show -p resource_dependencies_offline_restart asm-data1-rs

=== Resources ===

Resource: asm-data1-rs
Resource_dependencies_offline_restart: asm-inst-rs

— Standard and extension properties —
```

b. Set the new dependency.

- If `asm-data1-rs` already exists, use the following command to set the dependency. Note that the command includes the plus (+) symbol:

```
bash-3.00# clresource set -p resource_dependencies_offline_restart+=scal-asmdg1-rs
asm-data1-rs
```

- If `asm-data1-rs` does not yet exist, use the following command to create the resource with the offline-restart dependency:

```
bash-3.00# clresource create -g asm-dg-rg \
-t SUNW.scalable_asm_diskgroup_proxy \
-p asm_diskgroups=data1 \
-p resource_dependencies_offline_restart=asm-inst-rs,scal-asmdg1-rs \
-d asm-data1-rs
```

c. Verify the configured dependency.

```
bash-3.00# /u01/app/11.2.0/grid/bin/crsctl stat res ora.DATA1.dg -p | grep
START_DEPENDENCIES
START_DEPENDENCIES=hard(ora.asm,sun.scal-asmdg1-rs) pullup(ora.asm)
bash-3.00# clresource show -p resource_dependencies_offline_restart asm-data1-rs
=== Resources ===

Resource: asm-data1-rs
Resource_dependencies_offline_restart: asm-inst-rs scal-asmdg1-rs

— Standard and extension properties —
```

Top

Removing an Oracle Grid Infrastructure Resource

This section contains the following procedures to remove an Oracle Grid Infrastructure resource.

- How to Remove a Dependency
- How to Delete the `sun.`*`resource`* Resource

Top

How to Remove a Dependency

This procedure shows how to set the offline-restart dependency to remove a dependency.

1. Display the current dependency.

```
bash-3.00# /u01/app/11.2.0/grid/bin/crsctl stat res ora.DATA1.dg -p | grep START_DEPENDENCIES
START_DEPENDENCIES=hard(ora.asm,sun.scal-asmdg1-rs) pullup(ora.asm)
bash-3.00# clresource show -p resource_dependencies_offline_restart asm-data1-rs
=== Resources ===

Resource: asm-data1-rs
Resource_dependencies_offline_restart: asm-inst-rs scal-asmdg1-rs

— Standard and extension properties —
```

2. Set the new dependency.
   Note that the command includes the minus (-) symbol.

```
bash-3.00# clresource set -p resource_dependencies_offline_restart-=scal-asmdg1-rs asm-data1-rs
```

3. Verify the modified dependency.

```
bash-3.00# /u01/app/11.2.0/grid/bin/crsctl stat res ora.DATA1.dg -p | grep START_DEPENDENCIES
START_DEPENDENCIES=hard(ora.asm) pullup(ora.asm)
bash-3.00# clresource show -p resource_dependencies_offline_restart asm-data1-rs

=== Resources ===

Resource: asm-data1-rs
Resource_dependencies_offline_restart: asm-inst-rs

— Standard and extension properties —
```

Top

How to Delete the ***sun.resource*** Resource

1. Become superuser.
2. Ensure that the dependency is removed, as described in "How to Remove a Dependency" , and that `sun.resource` is stopped.

```
bash-3.00# /u01/app/11.2.0/grid/bin/crsctl stop res sun.scal-asmdg1-rs
CRS-2500: Cannot stop resource 'sun.scal-asmdg1-rs' as it is not running
CRS-4000: Command Stop failed, or completed with errors.
```

3. Delete `sun.resource`.

```
bash-3.00# /u01/app/11.2.0/grid/bin/crsctl delete res sun.scal-asmdg1-rs
```

4. Verify that `sun.resource` is deleted.

```
bash-3.00# /u01/app/11.2.0/grid/bin/crsctl stat res sun.scal-asmdg1-rs -p
CRS-210: Could not find resource 'sun.scal-asmdg1-rs'.
```

Top

Network-Attached Storage Manual

This section discusses errors, omissions, and additions in the Oracle Solaris Cluster 3.3 With Network-Attached Storage Device Manual.

Extraneous Information for Oracle's Sun Storage 7000 Unified Storage System as a NAS Device

At the initial publication of this manual, part number 821-1556-10, Chapter 3 documented the use of Oracle's Sun Storage 7000 Unified Storage System as a NAS device with fencing enabled. Support for Oracle's Sun Storage 7000 Unified Storage System as a NAS Device with fencing has not been enabled in the 3.3 release. Chapter 3 has subsequently been removed from the manual and republished as part number 821-1556-11. If you obtained a copy of this manual before Chapter 3 was removed (821-1556-10), ignore all information in Chapter 3 and information elsewhere in the manual that refers to this functionality.

Top

Sun NAS and Network Appliance NAS Directories Must be Exported to All Cluster Nodes (6996565)

For Sun NAS devices, you must explicitly grant access to each node in the cluster. All nodes in the cluster must have access to configured directories. When you add the Sun NAS directories to a cluster, ensure that the Sun NAS device is properly configured and the directories the cluster will use have been exported to all cluster nodes.

For Network Appliance NAS devices, directories must be exported to all cluster nodes. When you add the Network Appliance NAS directories to a cluster, ensure that the Network Appliance NAS device is properly configured and the directories the cluster will use have been exported to all cluster nodes.

Top

Man Pages

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

rt_properties(5)

The rt_properties (5) man page lists incorrect API versions. The correct API versions for the following Sun Cluster or Oracle Solaris Cluster releases are listed below.

Sun Cluster 3.2 1/09 API version 9
Sun Cluster 3.2 11/09 API version 10
Oracle Solaris Cluster 3.3 API version 11

Top

clnasdevice(1CL)

Some sections refer to the use of Oracle's Sun Storage 7000 Unified Storage System as a NAS device, using the NAS device type sun_uss. Support for Oracle's Sun Storage 7000 Unified Storage System as a NAS Device has not been enabled in the 3.3 release. Ignore all information for the {{sun_uss} type.

Top

clressharedaddress(1CL)

The clressharedaddress(1CL) man page is missing from the SUNWscman package. Refer instead to the online version of the man page:

http://download.oracle.com/docs/cd/E19680-01/821-1263/clressharedaddress-1cl?l=en&a=view

Top

scds_hasp_check(3HA)

The following information is missing from the DESCRIPTION:

Resource dependencies are only checked within the same cluster context in which the function is executed, either global cluster or zone cluster. Dependencies of the form *clustername*:*resourcename* (inter-cluster dependencies) are ignored. For example, if the only HAStoragePlus dependency is an inter-cluster dependency, the function returns the status code SCDS_HASP_NO_RESOURCE.

Top

SUNW.scalable_asm_diskgroup_proxy(5)

The SUNW.scalable_asm_diskgroup_proxy(5) is missing from the Oracle Solaris Cluster 3.3 release. The following are the man page contents.

Name

SUNW.scalable_asm_diskgroup_proxy, scalable_asm_diskgroup_proxy -- resource type implementation for the clustered Oracle Automated Storage Management (Oracle ASM) disk group managed by Oracle Solaris Cluster

Description

The SUNW.scalable_asm_diskgroup_proxy resource type represents the clustered Oracle ASM disk group in an Oracle Solaris Cluster configuration. This resource type is introduced in Oracle Solaris Cluster 3.3 for use starting with Oracle 11g release 2 configurations.

> **Note**
> The SUNW.scalable_asm_diskgroup_proxy resource type can only be used when using Oracle Grid Infrastructure for Clusters.

The SUNW.scalable_asm_diskgroup_proxy resource type is a multiple-master resource type. A single resource of this type can run on multiple nodes concurrently, but does not use network load balancing.

Each SUNW.scalable_asm_diskgroup_proxy resource represents a clustered Oracle ASM disk group. Each clustered Oracle ASM disk group is uniquely identified by the value of the asm_diskgroups extension property on the node where the instance is running. The Oracle ASM disk group resource should only be mounted if the Oracle ASM instance is available on the same cluster node. Additionally, Oracle ASM disk groups should only be mounted if any required storage resources are enabled on the cluster node. To ensure that these requirements are met, configure the Oracle ASM disk group resource as follows:

- Create a strong positive affinity between the Oracle ASM disk group resource group and the following resource groups:
  - The clustered Oracle ASM resource group
  - Any resource group that contains storage resources for Oracle files
- Create an offline-restart dependency between the clustered Oracle ASM disk group resource and the following resources:
  - The clustered Oracle ASM instance resource
  - Any storage resources for Oracle files that you are using

Create these dependencies and affinities when you configure clustered Oracle ASM disk group resources for the Oracle Solaris Cluster HA for Oracle data service or the Oracle Solaris Cluster Support for Oracle RAC data service. For more information about configuring resources for Oracle database instances, see Oracle Solaris Cluster Data Service for Oracle Guide or Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide.

To register this resource type and create instances of this resource type, use one of the following:

- Oracle Solaris Cluster Manager
- The clsetup(1CL) utility, specifying the option for configuring Oracle Solaris Cluster Support for Oracle Real Application Clusters
- The following sequence of Oracle Solaris Cluster maintenance commands:
  - To register this resource type, use the clresourcetype(1CL) command.
  - To create instances of this resource type, use the clresource(1CL) command.

Standard Properties

For a description of all standard resource properties, see the r_properties(5) man page.

Standard resource properties are overridden for this resource type as follows:

Prenet_start_timeout

Minimum	60

Default	300

`Prenet_stop_timeout`

Minimum	60
Default	300

Extension Properties

The extension properties of the `SUNW.scalable_asm_diskgroup_proxy` resource type are as follows.

`asm_diskgroups`

This property specifies the single-instance Oracle ASM disk group. If required, more than one single-instance Oracle ASM disk group can be specified as a comma separated list.

Data Type	String array
Range	Not applicable
Tunable	When disabled

`debug_level`

> **Note**
> All SQL*Plus and srvmgr messages that the Oracle ASM disk group resource issues are written to the log file `/var/opt/SUNWscor/oracle_asm/message_log.${RESOURCE}`.

This property indicates the level to which debug messages for the Oracle ASM disk-group resources are logged. When the debug level is increased, more debug messages are written to the system log `/var/adm/messages` as follows:

0	No debug messages
1	Function Begin and End messages
2	All debug messages and function Begin/End messages

Data Type	Integer
Range	0-2
Default	0
Tunable	Any time

`proxy_probe_timeout`

This property specifies the timeout value, in seconds, that the proxy monitor uses when checking the status of the clustered Oracle ASM disk group resource for which this resource is acting as a proxy.

Data Type	Integer
Range	5-120
Default	60
Tunable	Any time

`proxy_probe_interval`

This property specifies the interval, in seconds, between probes of the Oracle ASM disk group resource for which this resource is

acting as a proxy.

Data Type	Integer
Range	5-120
Default	30
Tunable	Any time

Example - Creating a ***scalable_asm_diskgroup_proxy*** Resource

This example shows the commands for performing the following operations to create a scalable multi-master `scalable_asm_diskgroup_proxy` resource on a two-node cluster:

- Creating the `asm-dg-rg` resource group
- Registering the `SUNW.scalable_asm_diskgroup_proxy` resource type
- Setting the resource group affinity
- Adding the `asm-dg-rs` resource to the `asm-dg-rg` resource group
- Setting the `asm_diskgroups` extension property for one Oracle ASM disk group

The example makes the following assumptions:

- The `bash` shell is used.
- A resource group that is named `asm-inst-rg` exists and contains a resource of type `SUNW.scalable_asm_instance_proxy` that is named `asm-inst-rs`.
- A resource group that is named `scal-mp-rg` exists and contains a resource of type `SUNW.ScalMountPoint` that is named `scal-mp-rs`, for Oracle files.

```
phys-schost-1# clresourcetype register SUNW.scalable_asm_diskgroup_proxy

phys-schost-1# clresourcegroup create -S asm-dg-rg
phys-schost-1# clresourcegroup set -p rg_affinities=++asm-inst-rg asm-dg-rg

phys-schost-1# clresource create -g asm-dg-rg \
-t SUNW.scalable_asm_diskgroup_proxy \
-p asm_diskgroups=data1 \
-p resource_dependencies_offline_restart=asm-inst-rs,qfs-mp-rs \
-d asm-dg-rs

phys-schost-1# clresourcegroup online -M asm-dg-rg

phys-schost-1# clresource enable asm-dg-rs
```

Attributes

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWscor

See Also

Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide, Oracle Solaris Cluster Data Service for Oracle Guide, Oracle Solaris Cluster Data Services Planning and Administration Guide

Top

`ORCL.obiee_cluster_controller`(5)

The `ORCL.obiee_cluster_controller`(5) man page is missing from the Oracle Solaris Cluster 3.3 release. The following are the man page contents.

Name

`ORCL.obiee_cluster_controller` -- resource type implementation for HA-Oracle Business Intelligence Enterprise Edition (Oracle BI EE) Cluster Controller.

Description

The `ORCL.obiee_cluster_controller` resource type represents the HA Oracle BI EE Cluster Controller in an Oracle Solaris Cluster configuration.

The HA-Oracle BI EE Cluster Controller resource is configured in a failover resource group when creating an Oracle BI EE installation that consists of more than one BI EE Server. The node lists for the resource groups that contain the primary and secondary cluster controllers must not contain common nodes.

You must set the following properties for an Oracle BI EE Cluster Controller resource by using the `clresource` command.

Standard Properties

The standard resource property `Failover` is set for all failover resource types.

See the `r_properties`(5) man page for a complete description of the following resource properties.

`Failover_mode`

Default:	SOFT
Tunable:	Any time

`Retry_count`

Default:	2
Tunable:	Any time

`Retry_interval`

Default:	1330
Tunable:	Any time

`Thorough_probe_interval`

Default:	30
Tunable:	Any time

Extension Properties

`BI_Install_Directory`

Type string (required). This property is set to the absolute path of the Oracle BI software installation directory.

You can modify this parameter only when disabled.

`BI_Data_Directory`

Type string (required). This property is set to the absolute path of the Oracle BI software data directory.

You can modify this parameter only when disabled.

`Run_64_bit`

Type Boolean (required). Default is TRUE. If this property is set to TRUE, then the 64-bit version of the binary is run. If set to FALSE, then the 32-bit version of the binary is run.

You can modify this parameter only when disabled.

`Cluster_Controller_Role`

Type enumerated (required). Default is PRIMARY. This property determines whether the resource represents a primary or secondary cluster controller. Valid values are PRIMARY or SECONDARY.

You can modify this parameter only when disabled.

`Debug_Level`

Type enumerated (optional). Default is 0. Valid values are 0, 1, 2. This property specifies the debug level.
You can modify this parameter at any time.

See Also

`pmfadm`(1M), `scha_resource_get`(1HA), `clresourcetype`(1CL), `clresource`(1CL)

Oracle Solaris Cluster Data Services Planning and Administration Guide

Top

`ORCL.obiee_presentation_service`(5)

The `ORCL.obiee_presentation_service`(5) man page is missing from the Oracle Solaris Cluster 3.3 release. The following are the man page contents.

Name

`ORCL.obiee_presentation_service` -- resource type implementation for HA-Oracle Business Intelligence Enterprise Edition (Oracle BI EE) Presentation Service.

Description

The `ORCL.obiee_presentation_service` resource type represents the HA-Oracle BI EE Presentation Service in an Oracle Solaris Cluster configuration. The HA-Oracle BI EE Presentation Service resource can be configured in one of two ways:

- In a failover resource group when the BI EE Cluster Controller is not used.
- In a multi-master resource group when the BI EE Cluster Controller is used.

You must set the following properties for an Oracle BI EE Presentation Service resource by using the `clresource` command.

Standard Properties

The standard resource property `Failover` is set for all failover resource types.

See `r_properties`(5) for a complete description of the following resource properties.

`Failover_mode`

Default:	SOFT
Tunable:	Any time

`Retry_count`

Default:	2
Tunable:	Any time

`Retry_interval`

Default:	1330

Tunable:	Any time

`Thorough_probe_interval`

Default:	30
Tunable:	Any time

Extension Properties

`BI_Install_Directory`

Type string (required). This property is set to the absolute path of the Oracle BI software installation directory.

You can modify this parameter only when disabled.

`BI_Data_Directory`

Type string (required). This property is set to the absolute path of the Oracle BI software data directory.

You can modify this parameter only when disabled.

`Run_64_bit`

Type Boolean (required). Default is TRUE. If this property is set to TRUE, then the 64-bit version of the binary is run. If set to FALSE, then the 32-bit version of the binary is run.

You can modify this parameter only when disabled.

`Debug_Level`

Type enumerated (optional). Default is 0. Valid values are 0, 1, 2. This property specifies the debug level.

You can modify this parameter at any time.

See Also

`pmfadm`(1M), `scha_resource_get`(1HA), `clresourcetype`(1CL), `clresource`(1CL)

Oracle Solaris Cluster Data Services Planning and Administration Guide

Top

`ORCL.obiee_scheduler`(5)

The `ORCL.obiee_scheduler`(5) man page is missing from the Oracle Solaris Cluster 3.3 release. The following are the man page contents.

Name

`ORCL.obiee_scheduler` -- resource type implementation for HA Oracle BI EE Scheduler.

Description

The `ORCL.obiee_scheduler` resource type represents the HA-Oracle Business Intelligence Enterprise Edition (Oracle BI EE) Scheduler in an Oracle Solaris Cluster configuration. The HA-Oracle BI EE Scheduler resource can be configured in one of two ways:

- In a failover resource group when the BI EE Cluster Controller is not used.
- In a multi-master resource group when the BI EE Cluster Controller is used.

You must set the following properties for an Oracle BI EE Scheduler resource by using the `clresource` command.

Standard Properties

The standard resource property `Failover` is set for all failover resource types.

See `r_properties`(5) for a complete description of the following resource properties.

`Failover_mode`

Default:	SOFT
Tunable:	Any time

`Retry_count`

Default:	2
Tunable:	Any time

`Retry_interval`

Default:	1330
Tunable:	Any time

`Thorough_probe_interval`

Default:	30
Tunable:	Any time

Extension Properties

`BI_Install_Directory`

Type string (required). This property is set to the absolute path of the Oracle BI software installation directory.

You can modify this parameter only when disabled.

`BI_Data_Directory`

Type string (required). This property is set to the absolute path of the Oracle BI software data directory.

You can modify this parameter only when disabled.

`Run_64_bit`

Type Boolean (required). Default is TRUE. If this property is set to TRUE, then the 64-bit version of the binary is run. If set to FALSE, then the 32-bit version of the binary is run.
You can modify this parameter only when disabled.

`Debug_Level`

Type enumerated (optional). Default is 0. Valid values are 0, 1, 2. This property specifies the debug level.
You can modify this parameter at any time.

See Also

`pmfadm`(1M), `scha_resource_get`(1HA), `clresourcetype`(1CL), `clresource`(1CL)

Oracle Solaris Cluster Data Services Planning and Administration Guide

Top

`ORCL.obiee_server`(5)

The ORCL.obiee_server(5) man page is missing from the Oracle Solaris Cluster 3.3 release. The following are the man page contents.

Name

ORCL.obiee_server -- resource type implementation for HA Oracle BI EE server.

Description

The ORCL.obiee_server resource type represents the HA-Oracle Business Intelligence Enterprise Edition (Oracle BI EE) Server in an Oracle Solaris Cluster configuration. The HA-Oracle BI EE Server resource can be configured in one of two ways:

- In a failover resource group when the BI EE Cluster Controller is not used.
- In a multi-master resource group when the BI EE Cluster Controller is used.

You must set the following properties for an Oracle BI EE Server resource by using the clresource command.

Standard Properties

The standard resource property Failover is set for all failover resource types.

See r_properties(5) for a complete description of the following resource properties.

Failover_mode

Default:	SOFT
Tunable:	Any time

Retry_count

Default:	2
Tunable:	Any time

Retry_interval

Default:	1330
Tunable:	Any time

Thorough_probe_interval

Default:	30
Tunable:	Any time

Extension Properties

BI_Install_Directory

Type string (required). This property is set to the absolute path of the Oracle BI software installation directory.
You can modify this parameter only when disabled.

BI_Data_Directory

Type string (required). This property is set to the absolute path of the Oracle BI software data directory.
You can modify this parameter only when disabled.

Run_64_bit

Type Boolean (required). Default is TRUE. If this property is set to TRUE, then the 64-bit version of the binary is run. If set to FALSE, then the 32-bit version of the binary is run.
You can modify this parameter only when disabled.

BI_PROBE_USER

Type string (optional). This property contains the Business Intelligence (BI) user name to use for probing the health of the BI server.
You can modify this parameter only when disabled.

BI_PROBE_PASSWORD

Type string (optional). This property contains the Business Intelligence (BI) user password to use for probing the health of the BI server.
You can modify this parameter only when disabled.

Debug_Level

Type enumerated (optional). Default is 0. Valid values are 0, 1, 2. This property specifies the debug level.
You can modify this parameter at any time.

See Also

`pmfadm`(1M), `scha_resource_get`(1HA), `clresourcetype`(1CL), `clresource`(1CL)

Oracle Solaris Cluster Data Services Planning and Administration Guide

Top

`SUNW.oracle_server`(5)

Oracle Solaris Cluster 3.3 introduces a new extension property for the `SUNW.oracle_server` resource type. This extension property, `db_unique_name`, is required when a single-instance HA for Oracle data service is deployed using Oracle 11g release 2 Grid Infrastructure for Clusters. The `db_unique_name` extension property specifies the unique name of the single-instance database that is being deployed.

When the `db_unique_name` extension property is used, the following extension properties are ignored and can be omitted from the HA-Oracle configuration:

```
alert_log_file
connect_string
```

The following example shows the creation of the `ora-db1-rg` resource group, where the unique database name `DB1` is specified for the `db_unique_name` extension property for the database. Because the `db_unique_name` extension property is used, the `connect_string` and `alert_log_file` extension properties are omitted.

```
# clrs create -g ora-db1-rg \
-t SUNW.oracle_server \
-p oracle_home=/u01/app/ora11g/product/11.2.0/dbhome_1 \
-p oracle_sid=DB1 \
-p db_unique_name=DB1 \
-p resource_dependencies_offline_restart=asm-dg1-rs \
-d ora-db1-rs
```

Top

Online Help

This section discusses errors, omissions, and additions in the online help.

Oracle Solaris Cluster Manager

The following information is missing from the online help for Oracle Solaris Cluster Manager:

- The pages that contain the Resource Group Status Properties table and the Zone Cluster Status Properties table are missing the following fields:
    - Priority - The order in which resource groups are assigned to master nodes. A higher priority indicates a more

important service. The default value is 500.
    - Preemption_Mode (Has_Cost/No_Cost/Never) - The likelihood that a resource group will be preempted from a node by a higher-priority resource group because of node overload. A Has_Cost value means that preempting this resource group has a cost associated with it. A No_Cost value indicates that the cost of preempting this resource group is zero. A Never value indicates that the resource group cannot be displaced from its current master to satisfy load limits.
    - Load Factors (Loadlimit@value) - How much of the load limit that the resource group consumes. The default value for each load factor is 0 and the maximum value is 1000.
- A new help page to describe how to create, edit, and delete a load limit:
    - Limitname - Name of the load limit.
    - Softlimit - The advisory upper boundary for a resource group load on a node or zone. The default value is 0, which indicates that no soft limit is imposed.
    - Hardlimit - Mandatory upper boundary for resource group load on a node or zone. The default value is null, which indicates that the "limit" is unlimited.

Top

Oracle BI EE Man Pages

ORCL.obiee_cluster_controller(5)

The `ORCL.obiee_cluster_controller`(5) man page is missing from the Oracle Solaris Cluster 3.3 release. The following are the man page contents.

Name

`ORCL.obiee_cluster_controller` -- resource type implementation for HA-Oracle Business Intelligence Enterprise Edition (Oracle BI EE) Cluster Controller.

Description

The `ORCL.obiee_cluster_controller` resource type represents the HA Oracle BI EE Cluster Controller in an Oracle Solaris Cluster configuration.

The HA-Oracle BI EE Cluster Controller resource is configured in a failover resource group when creating an Oracle BI EE installation that consists of more than one BI EE Server. The node lists for the resource groups that contain the primary and secondary cluster controllers must not contain common nodes.

You must set the following properties for an Oracle BI EE Cluster Controller resource by using the `clresource` command.

Standard Properties

The standard resource property `Failover` is set for all failover resource types.

See the `r_properties`(5) man page for a complete description of the following resource properties.

`Failover_mode`

Default:	SOFT
Tunable:	Any time

`Retry_count`

Default:	2
Tunable:	Any time

`Retry_interval`

Default:	1330
Tunable:	Any time

Thorough_probe_interval

Default:	30
Tunable:	Any time

Extension Properties

BI_Install_Directory

Type string (required). This property is set to the absolute path of the Oracle BI software installation directory.

You can modify this parameter only when disabled.

BI_Data_Directory

Type string (required). This property is set to the absolute path of the Oracle BI software data directory.

You can modify this parameter only when disabled.

Run_64_bit

Type Boolean (required). Default is TRUE. If this property is set to TRUE, then the 64-bit version of the binary is run. If set to FALSE, then the 32-bit version of the binary is run.

You can modify this parameter only when disabled.

Cluster_Controller_Role

Type enumerated (required). Default is PRIMARY. This property determines whether the resource represents a primary or secondary cluster controller. Valid values are PRIMARY or SECONDARY.

You can modify this parameter only when disabled.

Debug_Level

Type enumerated (optional). Default is 0. Valid values are 0, 1, 2. This property specifies the debug level.
You can modify this parameter at any time.

See Also

`pmfadm`(1M), `scha_resource_get`(1HA), `clresourcetype`(1CL), `clresource`(1CL)

Oracle Solaris Cluster Data Services Planning and Administration Guide

Top

ORCL.obiee_presentation_service(5)

The `ORCL.obiee_presentation_service`(5) man page is missing from the Oracle Solaris Cluster 3.3 release. The following are the man page contents.

Name

`ORCL.obiee_presentation_service` -- resource type implementation for HA-Oracle Business Intelligence Enterprise Edition (Oracle BI EE) Presentation Service.

Description

The `ORCL.obiee_presentation_service` resource type represents the HA-Oracle BI EE Presentation Service in an Oracle Solaris Cluster configuration. The HA-Oracle BI EE Presentation Service resource can be configured in one of two ways:

- In a failover resource group when the BI EE Cluster Controller is not used.
- In a multi-master resource group when the BI EE Cluster Controller is used.

You must set the following properties for an Oracle BI EE Presentation Service resource by using the `clresource` command.

Standard Properties

The standard resource property `Failover` is set for all failover resource types.

See `r_properties`(5) for a complete description of the following resource properties.

`Failover_mode`

Default:	SOFT
Tunable:	Any time

`Retry_count`

Default:	2
Tunable:	Any time

`Retry_interval`

Default:	1330
Tunable:	Any time

`Thorough_probe_interval`

Default:	30
Tunable:	Any time

Extension Properties

`BI_Install_Directory`

Type string (required). This property is set to the absolute path of the Oracle BI software installation directory.

You can modify this parameter only when disabled.

`BI_Data_Directory`

Type string (required). This property is set to the absolute path of the Oracle BI software data directory.

You can modify this parameter only when disabled.

`Run_64_bit`

Type Boolean (required). Default is TRUE. If this property is set to TRUE, then the 64-bit version of the binary is run. If set to FALSE, then the 32-bit version of the binary is run.

You can modify this parameter only when disabled.

`Debug_Level`

Type enumerated (optional). Default is 0. Valid values are 0, 1, 2. This property specifies the debug level.

You can modify this parameter at any time.

See Also

`pmfadm`(1M), `scha_resource_get`(1HA), `clresourcetype`(1CL), `clresource`(1CL)

Oracle Solaris Cluster Data Services Planning and Administration Guide

Top

ORCL.obiee_scheduler(5)

The ORCL.obiee_scheduler(5) man page is missing from the Oracle Solaris Cluster 3.3 release. The following are the man page contents.

Name

ORCL.obiee_scheduler -- resource type implementation for HA Oracle BI EE Scheduler.

Description

The ORCL.obiee_scheduler resource type represents the HA-Oracle Business Intelligence Enterprise Edition (Oracle BI EE) Scheduler in an Oracle Solaris Cluster configuration. The HA-Oracle BI EE Scheduler resource can be configured in one of two ways:

- In a failover resource group when the BI EE Cluster Controller is not used.
- In a multi-master resource group when the BI EE Cluster Controller is used.

You must set the following properties for an Oracle BI EE Scheduler resource by using the clresource command.

Standard Properties

The standard resource property Failover is set for all failover resource types.

See r_properties(5) for a complete description of the following resource properties.

Failover_mode

Default:	SOFT
Tunable:	Any time

Retry_count

Default:	2
Tunable:	Any time

Retry_interval

Default:	1330
Tunable:	Any time

Thorough_probe_interval

Default:	30
Tunable:	Any time

Extension Properties

BI_Install_Directory

Type string (required). This property is set to the absolute path of the Oracle BI software installation directory.

You can modify this parameter only when disabled.

BI_Data_Directory

Type string (required). This property is set to the absolute path of the Oracle BI software data directory.

You can modify this parameter only when disabled.

Run_64_bit

Type Boolean (required). Default is TRUE. If this property is set to TRUE, then the 64-bit version of the binary is run. If set to FALSE, then the 32-bit version of the binary is run.
You can modify this parameter only when disabled.

Debug_Level

Type enumerated (optional). Default is 0. Valid values are 0, 1, 2. This property specifies the debug level.
You can modify this parameter at any time.

See Also

`pmfadm`(1M), `scha_resource_get`(1HA), `clresourcetype`(1CL), `clresource`(1CL)

Oracle Solaris Cluster Data Services Planning and Administration Guide

Top

ORCL.obiee_server(5)

The `ORCL.obiee_server`(5) man page is missing from the Oracle Solaris Cluster 3.3 release. The following are the man page contents.

Name

`ORCL.obiee_server` -- resource type implementation for HA Oracle BI EE server.

Description

The `ORCL.obiee_server` resource type represents the HA-Oracle Business Intelligence Enterprise Edition (Oracle BI EE) Server in an Oracle Solaris Cluster configuration. The HA-Oracle BI EE Server resource can be configured in one of two ways:

- In a failover resource group when the BI EE Cluster Controller is not used.
- In a multi-master resource group when the BI EE Cluster Controller is used.

You must set the following properties for an Oracle BI EE Server resource by using the `clresource` command.

Standard Properties

The standard resource property `Failover` is set for all failover resource types.

See `r_properties`(5) for a complete description of the following resource properties.

Failover_mode

Default:	SOFT
Tunable:	Any time

Retry_count

Default:	2
Tunable:	Any time

Retry_interval

Default:	1330
Tunable:	Any time

Thorough_probe_interval

Default:	30

Tunable:	Any time

Extension Properties

BI_Install_Directory

Type string (required). This property is set to the absolute path of the Oracle BI software installation directory.
You can modify this parameter only when disabled.

BI_Data_Directory

Type string (required). This property is set to the absolute path of the Oracle BI software data directory.
You can modify this parameter only when disabled.

Run_64_bit

Type Boolean (required). Default is TRUE. If this property is set to TRUE, then the 64-bit version of the binary is run. If set to FALSE, then the 32-bit version of the binary is run.
You can modify this parameter only when disabled.

BI_PROBE_USER

Type string (optional). This property contains the Business Intelligence (BI) user name to use for probing the health of the BI server.
You can modify this parameter only when disabled.

BI_PROBE_PASSWORD

Type string (optional). This property contains the Business Intelligence (BI) user password to use for probing the health of the BI server.
You can modify this parameter only when disabled.

Debug_Level

Type enumerated (optional). Default is 0. Valid values are 0, 1, 2. This property specifies the debug level.
You can modify this parameter at any time.

See Also

`pmfadm`(1M), `scha_resource_get`(1HA), `clresourcetype`(1CL), `clresource`(1CL)

Oracle Solaris Cluster Data Services Planning and Administration Guide

Top

Changes With Apache Tomcat 5.5.28 and 6.0.29

Changes With Apache Tomcat 5.5.28 and 6.0.29

The following information additions and changes support the `wget` probe algorithm for Oracle Solaris Cluster HA for Apache Tomcat (HA for Apache Tomcat) using Apache Tomcat versions starting with 5.5.28 and 6.0.29. This functionality is delivered in the following Oracle Solaris Cluster 3.3 patches for HA for Apache Tomcat:

- SPARC: 145341-01
- x86: 145342-01

New Restriction For Apache Tomcat 5.5.28 and 6.0.29

Starting with Apache Tomcat versions 5.5.28 and 6.0.29, you must use the `wget` probe algorithm. This bypasses the normal probe, enabling the TestUrl parameter to work correctly. You can optionally use the `wget` probe algorithm for earlier versions as well.

Top

New Requirement for Apache Tomcat 5.5.28 and 6.0.29

To support the `wget` probe algorithm, the `wget` command must be available from every node of the cluster.

Top

Installing and Configuring Apache Tomcat

In the bullet list of considerations to determine how Apache Tomcat will be deployed with Oracle Solaris Cluster, you must also determine which Apache Tomcat version to deploy.

Top

How to Register and Configure Solaris Cluster HA for Apache Tomcat as a Failover Data Service

The following replaces Step 1:

1. Prepare the parameter file, which is required by HA for Apache Tomcat.

```
# cd /opt/SUNWsctomcat/bin
# cp pfile desired-place
```

Edit the parameter file `pfile` and follow the comments within that file. For example:

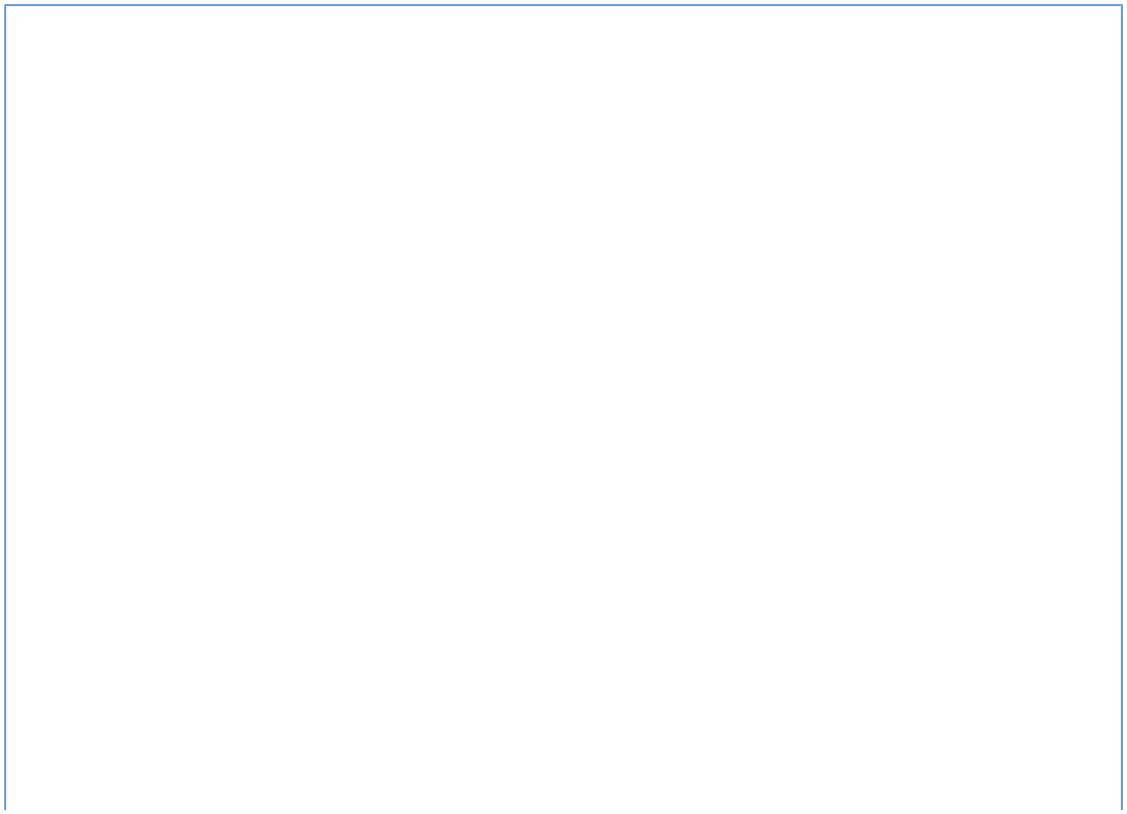

```
#!/usr/bin/ksh
#
# CDDL HEADER START
#
# The contents of this file are subject to the terms of the
# Common Development and Distribution License (the License).
# You may not use this file except in compliance with the License.
#
# You can obtain a copy of the license at usr/src/CDDL.txt
# or http://www.opensolaris.org/os/licensing.
# See the License for the specific language governing permissions
# and limitations under the License.
#
# When distributing Covered Code, include this CDDL HEADER in each
# file and include the License file at usr/src/CDDL.txt.
# If applicable, add the following below this CDDL HEADER, with the
# fields enclosed by brackets \[\] replaced with your own identifying
# information: Portions Copyright \[yyyy\] \[name of copyright owner\]
#
# CDDL HEADER END
#
#
# Copyright (c) 2006, 2010, Oracle and/or its affiliates. All rights reserved.
#

#ident  "@(#)pfile.ksh  1.6     10/08/16"

# Set the Apache Tomcat specific environment variables which the start, stop
# and check fuctions will use
#
# EnvScript     Script to set runtime environment for tomcat
# User          Apache Tomcat User
# Basepath      Absolute path to Tomcat Home directory i.e. 4.x CATALINA_HOME
#               or TOMCAT_HOME for 3.x
# Host          Hostname to test Apache Tomcat
# Port          Port where Apache Tomcat is configured
# TestCmd       Apache Tomcat test command
#
# Some versions of Apache Tomcat do not react on the probe algorithm with the
# http get command. Because of this we provide an alternative probe method
# which uses wget. If wget is not installed on your system you must install it
# manually and provide the WgetPath variable. If the WgetPath variable is not set,
# it defaults to /usr/sfw/bin/wget.
# You can use the wget algorithm for https probe as well which are not possible
# with the http get command.
#
# Use the variables TestUrl and WgetPath as a replacement for:
# Port, Host, TestCmd
# You can specify only one set of variables.
#
# TestUrl       Url where Tomcat serve a web site. This can be done either
#               by the http or https  protocol
#               Format: http://hostname:port/startpage
#               Example: http://localhost:8080/
#               You might want to include options here like:
#               "--no-cookies http://localhost:8080/"
#               The requirement here is that wget -O - ${TestUrl} produces
#               the output containing the ReturnString.
# WgetPath      Abolute pat to wget, leaving this variable empty defaults to:
#               /usr/sfw/bin/wget
# Startwait     Sleeping $Startwait seconds after completion of the
#               start command
# ReturnString  This string must be present in the output of the http get command
#               or in the output of the wget command.
EnvScript=
User=
Basepath=
Host=
Port=8080
TestCmd="get /index.jsp"
TestUrl=
WgetPath=
ReturnString="CATALINA"
Startwait=20
```

The following is an example for Apache Tomcat 4.1.24:

```
EnvScript=/global/mnt1/tomtest/env.ksh
User=tomtest
Basepath=/global/mnt1/tomtest/jakarta-tomcat-4.1.24
Host=tomcat-lh
Port=8080
TestCmd="get /index.jsp"
TestUrl=
WgetPath=
ReturnString="CATALINA"
Startwait=20
```

This example for Apache Tomcat 4.1.24 configures a standalone web server with the default start page `index.jsp`. The start page contains the string `CATALINA`.

To use the alternative `wget` probe, assuming `wget` is available under `/usr/sfw/bin/wget`, instead set the following variables:

```
EnvScript=/global/mnt1/tomtest/env.ksh
User=tomtest
Basepath=/global/mnt1/tomtest/jakarta-tomcat-4.1.24
Host=
Port=
TestCmd=
TestUrl=”http://tomcat-lh:8080/”
WgetPath=
ReturnString="CATALINA"
Startwait=20
```

Note
The parameter files must be available on every node that can host the Apache Tomcat data service. For a scalable or multiple-masters configuration, store the parameter files on the local file system of every node or non-global zone or on the shared storage. The parameter files must not differ for any instance of Apache Tomcat on the various nodes.

Repeat Step 1 for every Apache Tomcat instance you need.

Top

How to Register and Configure Solaris Cluster HA for Apache Tomcat as a Multiple Masters Data Service

The following replaces Step 2:

2. Prepare the parameter file, which is required by HA for Apache Tomcat.

```
# cd /opt/SUNWsctomcat/util
# cp sctomcat_config desired-place
```

Edit the parameter file `pfile` and follow the comments within that file. For example:

```
#!/usr/bin/ksh
#
# CDDL HEADER START
#
# The contents of this file are subject to the terms of the
# Common Development and Distribution License (the License).
# You may not use this file except in compliance with the License.
#
# You can obtain a copy of the license at usr/src/CDDL.txt
# or http://www.opensolaris.org/os/licensing.
# See the License for the specific language governing permissions
# and limitations under the License.
#
```

```
# When distributing Covered Code, include this CDDL HEADER in each
# file and include the License file at usr/src/CDDL.txt.
# If applicable, add the following below this CDDL HEADER, with the
# fields enclosed by brackets \[\] replaced with your own identifying
# information: Portions Copyright \[yyyy\] \[name of copyright owner\]
#
# CDDL HEADER END
#

#
# Copyright (c) 2006, 2010, Oracle and/or its affiliates. All rights reserved.
#

#ident  "@(#)pfile.ksh  1.6     10/08/16"

# Set the Apache Tomcat specific environment variables which the start, stop
# and check fuctions will use
#
# EnvScript     Script to set runtime environment for tomcat
# User          Apache Tomcat User
# Basepath      Absolute path to Tomcat Home directory i.e. 4.x CATALINA_HOME
#               or TOMCAT_HOME for 3.x
# Host          Hostname to test Apache Tomcat
# Port          Port where Apache Tomcat is configured
# TestCmd       Apache Tomcat test command
#
# Some versions of Apache Tomcat do not react on the probe algorithm with the
# http get command. Because of this we provide an alternative probe method
# which uses wget. If wget is not installed on your system you must install it
# manually and provide the WgetPath variable. If the WgetPath variable is not set,
# it defaults to /usr/sfw/bin/wget.
# You can use the wget algorithm for https probe as well which are not possible
# with the http get command.
#
# Use the variables TestUrl and WgetPath as a replacement for:
# Port, Host, TestCmd
# You can specify only one set of variables.
#
# TestUrl       Url where Tomcat serve a web site. This can be done either
#               by the http or https  protocol
#               Format: http://hostname:port/startpage
#               Example: http://localhost:8080/
#               You might want to include options here like:
#               "--no-cookies http://localhost:8080/"
#               The requirement here is that wget -O - ${TestUrl} produces
#               the output containing the ReturnString.
# WgetPath      Abolute pat to wget, leaving this variable empty defaults to:
#               /usr/sfw/bin/wget
# Startwait     Sleeping $Startwait seconds after completion of the
#               start command
# ReturnString  This string must be present in the output of the http get command
#               or in the output of the wget command.
EnvScript=
User=
Basepath=
Host=
Port=8080
TestCmd="get /index.jsp"
TestUrl=
WgetPath=
```

```
ReturnString="CATALINA"
Startwait=20
```

The following is an example for Apache Tomcat 4.1.24.

```
EnvScript=/tomcat/env.ksh
User=tomcat
Basepath=/tomcat/jakarta-tomca-4.1.24
Host=localhost
Port=8080
TestCmd="get /index.jsp"
TestUrl=
WgetPath=
ReturnString="CATALINA"
Startwait=20
```

This example for Apache Tomcat 4.1.24 configures a standalone web server with the default start page `index.jsp`. The start page contains the string `CATALINA`.

Depending on the selected probe method, the Host parameter or the IP alias in the TestUrl parameter depends on the location of the `pfile`:

- If the `pfile` is stored on the shared storage, you must use `localhost`.
- If the `pfile` is stored on the local storage, it can be either `localhost` or the node's hostname.

To use the alternative `wget` probe, assuming `wget` is available under `/usr/sfw/bin/wget`, set the following variables:

```
EnvScript=/global/mnt1/tomtest/env.ksh
User=tomtest
Basepath=/global/mnt1/tomtest/jakarta-tomcat-4.1.24
Host=
Port=
TestCmd=
TestUrl=”http://localhost:8080/”
WgetPath=
ReturnString="CATALINA"
Startwait=20
```

Note
The parameter files must be available on every node that can host the Apache Tomcat data service. For a scalable or multiple-masters configuration, store the parameter files on the local file system of every node or non-global zone or on the shared storage. The parameter files must not differ for any instance of Apache Tomcat on the various nodes.

Repeat this Step for every Apache Tomcat instance and every node you need.

Top

How to Register and Configure Solaris Cluster HA for Apache Tomcat as a Scalable Data Service

The following replaces Step 2:

2. Prepare the parameter file, which is required by Oracle Solaris Cluster HA for Apache Tomcat.

```
# cd /opt/SUNWsctomcat/bin
# cp pfile desired-place
```

Edit the parameter file `pfile` and follow the comments within that file. For example:

```
#!/usr/bin/ksh
#
# CDDL HEADER START
```

```
#
# The contents of this file are subject to the terms of the
# Common Development and Distribution License (the License).
# You may not use this file except in compliance with the License.
#
# You can obtain a copy of the license at usr/src/CDDL.txt
# or http://www.opensolaris.org/os/licensing.
# See the License for the specific language governing permissions
# and limitations under the License.
#
# When distributing Covered Code, include this CDDL HEADER in each
# file and include the License file at usr/src/CDDL.txt.
# If applicable, add the following below this CDDL HEADER, with the
# fields enclosed by brackets \[\] replaced with your own identifying
# information: Portions Copyright \[yyyy\] \[name of copyright owner\]
#
# CDDL HEADER END
#

#
# Copyright (c) 2006, 2010, Oracle and/or its affiliates. All rights reserved.
#

#ident  "@(#)pfile.ksh  1.6     10/08/16"

# Set the Apache Tomcat specific environment variables which the start, stop
# and check fuctions will use
#
# EnvScript     Script to set runtime environment for tomcat
# User          Apache Tomcat User
# Basepath      Absolute path to Tomcat Home directory i.e. 4.x CATALINA_HOME
#               or TOMCAT_HOME for 3.x
# Host          Hostname to test Apache Tomcat
# Port          Port where Apache Tomcat is configured
# TestCmd       Apache Tomcat test command
#
# Some versions of Apache Tomcat do not react on the probe algorithm with the
# http get command. Because of this we provide an alternative probe method
# which uses wget. If wget is not installed on your system you must install it
# manually and provide the WgetPath variable. If the WgetPath variable is not set,
# it defaults to /usr/sfw/bin/wget.
# You can use the wget algorithm for https probe as well which are not possible
# with the http get command.
#
# Use the variables TestUrl and WgetPath as a replacement for:
# Port, Host, TestCmd
# You can specify only one set of variables.
#
# TestUrl       Url where Tomcat serve a web site. This can be done either
#               by the http or https  protocol
#               Format: http://hostname:port/startpage
#               Example: http://localhost:8080/
#               You might want to include options here like:
#               "--no-cookies http://localhost:8080/"
#               The requirement here is that wget -O - ${TestUrl} produces
#               the output containing the ReturnString.
# WgetPath      Abolute pat to wget, leaving this variable empty defaults to:
#               /usr/sfw/bin/wget
# Startwait     Sleeping $Startwait seconds after completion of the
#               start command
# ReturnString  This string must be present in the output of the http get command
#               or in the output of the wget command.
EnvScript=
User=
Basepath=
Host=
Port=8080
TestCmd="get /index.jsp"
TestUrl=
WgetPath=
```

```
ReturnString="CATALINA"
Startwait=20
```

The following is an example for Apache Tomcat 4.1.24.

```
EnvScript=/tomcat/env.ksh
User=tomcat
Basepath=/tomcat/jakarta-tomca-4.1.24
Host=localhost
Port=8080
TestCmd="get /index.jsp"
TestUrl=
WgetPath=
ReturnString="CATALINA"
Startwait=20
```

To use the alternative `wget` probe, assuming `wget` is available under `/usr/sfw/bin/wget`, set the following variables:

```
EnvScript=/global/mnt1/tomtest/env.ksh
User=tomtest
Basepath=/global/mnt1/tomtest/jakarta-tomcat-4.1.24
Host=
Port=
TestCmd=
TestUrl=”http://localhost:8080/”
WgetPath=
ReturnString="CATALINA"
Startwait=20
```

This example for Apache Tomcat 4.1.24 configures a standalone web server with the default start page `index.jsp`. The start page contains the string `CATALINA`.

Depending on the selected probe method, the Host parameter or the IP alias in the TestUrl parameter depends on the location of the `pfile`.

- If the `pfile` is stored on the shared storage, you must use `localhost`.
- If the `pfile` is stored on the local storage, it can be either `localhost` or the node's hostname.

Note
Do not use the SharedAddress here.

Note
The parameter files must be available on every node that can host the Apache Tomcat data service. For a scalable or multiple-masters configuration, store the parameter files on the local file system of every node or on the shared storage. The parameter files must not differ for an instance of Apache Tomcat on the various nodes.

Repeat this Step for every Apache Tomcat instance and every node you need.

Top

Structure of the Apache Tomcat HA Parameter File

The following are additions to Table 6: Structure of the Apache Tomcat HA parameter file.

Variable	Explanation
TestUrl	This variable presents a complete `http` or `https` URL, which is passed to the `wget` command to test the Tomcat server's sanity.
WgetPath	This variable presents the absolute path to the `wget` binary, if the `wget` binary is not located under `/usr/sfw/bin/wget`.

In addition, TestUrl and WgetPath are added to the list of parameters that are used for starting and stopping Apache Tomcat.

Top

Strategy to Choose the TestCmd and the ReturnString Variable

The following replaces the section "Strategy to Chose the TestCmd and the ReturnString Variable".

Strategy to Choose the TestCmd or TestUrl Variable and the ReturnString Variable

The following alternatives are available to choose the ReturnString variable:

- Take the start page of your application and set the TestCmd to get */start_page*. If you picked the `wget` probe method, set TestUrl to `http://`*`ipalias`*`:`*`port`*`/start_page`. With this strategy, you are monitoring that Apache Tomcat is serving your application.
- Take the Apache Tomcat default start page and set the TestCmd to get */default-startpage*. If you picked the `wget` method, set TestUrl to `http://`*`ipalias`*`:`*`port`*`/`*`default-startpage`*. In this case, set the ReturnString to a string contained in the startpage. This string depends on the deployed Apache Tomcat version; for 3.x it is Tomcat home and for 4.x and 5.x it is CATALINA. With this strategy, you are monitoring that Apache Tomcat is serving its default application.
- Deploy a test application (which is not provided with the HA for Apache Tomcat agent) to Apache Tomcat. Set the TestCmd to get *`startpage_of_the_application`*. If you use the `wget` method, set TestUrl to `http://`*`ipalias`*`:`*`port`*`/startpage_of_the_application`. In this case, set the ReturnString to a string contained in the start page. With this strategy, you are monitoring that Apache Tomcat is serving your test application.
- If none of above is appropriate, set the TestCmd to get */a-page-that-does-not-exist*. In this case, set the ReturnString to a string contained in the Error Page. With this strategy, you are monitoring that Apache Tomcat is operational, because it registers that it must deliver a page that does not exist.

You can evaluate the different pages by connecting from a browser to *`hostname`*`:`*`port`* and specifying the different pages.

Top

Probing Algorithm and Functionality

In the steps that describe how the sanity of Apache Tomcat is monitored, the following replaces Step 2:

2. Pings the Host, which is configured in the Apache Tomcat HA parameter file, unless the TestUrl parameter is specified in the Tomcat parameter file.

Top

Deployment Example: Installing Apache Tomcat in the Global Zone

In the section "Software Configuration", the following is an additional expectation of the example configuration:

- The `wget` binary is available under `/usr/sfw/bin/wget`.

Top

Example: Modifying the Apache Tomcat Configuration

The following is an addition to Step 4:

As an alternative, you can use the following set of variables:

```
EnvScript=/global/mnt3/tomcat/env.ksh
User=tomcat
Basepath=/global/mnt3/tomcat/apache-tomcat-5.5.17
TestUrl="http://ha-host-1:8080/"
ReturnString="CATALINA"
Startwait=20
```

Top

Deployment Example: Installing Apache Tomcat in a Failover Zone

In the section "Software Configuration", the following is an additional expectation of the example configuration:

- The `wget` binary is available under `/usr/sfw/bin/wget`.

Top

Example: Installing the Apache Tomcat Software on Shared Storage

The following is an addition to Step 6:

As an alternative, you can use the following set of variables:

```
EnvScript=/global/mnt3/tomcat/env.ksh
User=tomcat
Basepath=/global/mnt3/tomcat/apache-tomcat-5.5.17
TestUrl="http://ha-host-1:8080/"
ReturnString="CATALINA"
Startwait=20
```

Top

Deployment Example: Installing Apache Tomcat in a Non-Global Zone

In the section "Software Configuration", the following is an additional expectation of the example configuration:

- The `wget` binary is available under `/usr/sfw/bin/wget`.

Top

Example: Modifying the Apache Tomcat Configuration File

The following is an addition to Step 3:

As an alternative, you can use the following set of variables:

```
EnvScript=/global/mnt3/tomcat/env.ksh
User=tomcat
Basepath=/global/mnt3/tomcat/apache-tomcat-5.5.17
TestUrl="http://ha-host-1:8080/"
ReturnString="CATALINA"
Startwait=20
```

Top