Oracle® Solaris Cluster Data Service for Informix Guide

ORACLE®

Contents

Preface

Oracle Solaris Cluster Data Service for Informix Guide explains how to install and configure HA for Informix.

Note – This Oracle Solaris Cluster release supports systems that use the SPARC and x86 families of processor architectures: UltraSPARC, SPARC64, AMD64, and Intel 64. In this document, x86 refers to the larger family of 64-bit x86 compatible products. Information in this document pertains to all platforms unless otherwise specified.

This document is intended for system administrators with extensive knowledge of Oracle software and hardware. Do not use this document as a planning or presales guide. Before reading this document, you should have already determined your system requirements and purchased the appropriate equipment and software.

The instructions in this book assume knowledge of the Solaris Operating System (Solaris OS) and expertise with the volume-manager software that is used with Oracle Solaris Cluster software.

Using UNIX Commands

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

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

Typographic Conventions

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

TABLE P–1 Typographic Conventions

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

Shell Prompts in Command Examples

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

TABLE P–2 Shell Prompts

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

Related Documentation

Information about related Oracle Solaris Cluster topics is available in the documentation that is listed in the following table. All Oracle Solaris Cluster documentation is available at `http://docs.sun.com`.

Topic	Documentation
Data service administration	*Oracle Solaris Cluster Data Services Planning and Administration Guide*
	Individual data service guides
Concepts	*Oracle Solaris Cluster Concepts Guide*
Overview	*Oracle Solaris Cluster Overview*
Software installation	*Oracle Solaris Cluster Software Installation Guide*
System administration	*Oracle Solaris Cluster System Administration Guide*
Hardware administration	*Oracle Solaris Cluster 3.3 Hardware Administration Manual*
	Individual hardware administration guides
Data service development	*Oracle Solaris Cluster Data Services Developer's Guide*
Error messages	*Oracle Solaris Cluster Error Messages Guide*
Command and function reference	*Oracle Solaris Cluster Reference Manual*

For a complete list of Oracle Solaris Cluster documentation, see the release notes for your release of Oracle Solaris Cluster at `http://docs.sun.com`.

Related Third-Party Web Site References

Third-party URLs that are referenced in this document provide additional related information.

Note – Oracle is not responsible for the availability of third-party web sites mentioned in this document. Oracle does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Oracle will not be responsible or liable for any actual or alleged damage or loss caused or alleged to be caused by or in connection with use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

Documentation, Support, and Training

See the following web sites for additional resources:

- Documentation (http://docs.sun.com)
- Support (http://www.oracle.com/us/support/systems/index.html)
- Training (http://education.oracle.com) – Click the Sun link in the left navigation bar.

Oracle Welcomes Your Comments

Oracle welcomes your comments and suggestions on the quality and usefulness of its documentation. If you find any errors or have any other suggestions for improvement, go to http://docs.sun.com and click Feedback. Indicate the title and part number of the documentation along with the chapter, section, and page number, if available. Please let us know if you want a reply.

Oracle Technology Network (http://www.oracle.com/technetwork/index.html) offers a range of resources related to Oracle software:

- Discuss technical problems and solutions on the Discussion Forums (http://forums.oracle.com).
- Get hands-on step-by-step tutorials with Oracle By Example (http://www.oracle.com/technology/obe/start/index.html).
- Download Sample Code (http://www.oracle.com/technology/sample_code/index.html).

Getting Help

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

- Your name and email address (if available)

- Your company name, address, and phone number

- The model number and serial number of your systems

- The release number of the Oracle Solaris Operating System (for example, Oracle Solaris 10)

- The release number of Oracle Solaris Cluster (for example, Oracle Solaris Cluster 3.3)

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

Command	Function
prtconf -v	Displays the size of the system memory and reports information about peripheral devices
psrinfo -v	Displays information about processors
showrev –p	Reports which patches are installed
prtdiag -v	Displays system diagnostic information
/usr/cluster/bin/clnode show-rev	Displays Oracle Solaris Cluster release and package version information

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

◆ ◆ ◆

1

Installing and Configuring Oracle Solaris Cluster HA for Informix

This chapter explains how to install and configure the Oracle Solaris Cluster HA for Informix data service.

This chapter contains the following sections:

Overview of the Installation and Configuration Process for Oracle Solaris Cluster HA for Informix

The table below lists the tasks for installing and configuring Oracle Solaris Cluster HA for Informix. Perform these tasks in the order that they are listed.

TABLE 1–1   Installing and Configuring Oracle Solaris Cluster HA for Informix

Task	Instruction

TABLE 1–1 Installing and Configuring Oracle Solaris Cluster HA for Informix   *(Continued)*

Plan the Informix installation	"Planning the Oracle Solaris Cluster HA for Informix Installation and Configuration" on page 12
Prepare the nodes and disks for Informix installation	"Preparing the Nodes and Disks" on page 14
Install and configure Informix	"Installing and Configuring Informix" on page 14
Enable the Informix instances to run in a cluster	"Enabling the Informix Instances to Run in a Cluster" on page 15
Install Oracle Solaris Cluster HA for Informix packages	"Installing the HA for Informix Packages" on page 17
Register and configure Oracle Solaris Cluster HA for Informix	"Registering and Configuring Oracle Solaris Cluster HA for Informix" on page 19
Verify Oracle Solaris Cluster HA for Informix installation and configuration	"Verifying the Oracle Solaris Cluster HA for Informix Installation and Configuration" on page 19
Understand the Oracle Solaris Cluster HA for Informix Fault monitor	"Understanding the Oracle Solaris Cluster HA for Informix Fault Monitor" on page 20
Debug Oracle Solaris Cluster HA for Informix	"Debugging Oracle Solaris Cluster HA for Informix" on page 21

Planning the Oracle Solaris Cluster HA for Informix Installation and Configuration

This section contains the information you need to plan your Oracle Solaris Cluster HA for Informix installation and configuration.

Note – Oracle Solaris Cluster HA for Informix can be configured to run in a whole root or a sparse root non-global zone, if required.

Configuration Restrictions

This section provides a list of software and hardware configuration restrictions that apply to Oracle Solaris Cluster HA for Informix. Use the restrictions in this section to plan the installation and configuration of Oracle Solaris Cluster HA for Informix.

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

For restrictions that apply to all data services, see the *Sun Cluster Release Notes for Solaris OS*.

Restriction for Supported Configurations of Oracle Solaris Cluster HA for Informix

Oracle Solaris Cluster HA for Informix data service can be configured only as a failover service. Single or multiple instances of Informix Dynamic Server can be deployed in the cluster. Informix Dynamic Server can be deployed in the global zone or non-global zone.

Restriction for the Location of Informix Files

You need to keep these restrictions in mind while planning the installation and configuration of Oracle Solaris Cluster HA for Informix.

- **Informix binary files** – You need to place the INFORMIXDIR directory on a shared storage such as a cluster file system or a highly available local file system.

  Alternatively, you can install Informix locally on each node. In this case, the INFORMIXDIR directory will reflect a local file system.

  If you decide to install Informix locally on each node, you must ensure that the ONCONFIG and INFORMIXSQLHOSTS files are kept in synchronization across the cluster nodes.

- **Informix database files** – You must ensure that the Informix database files are placed on a shared storage. You can use a highly available local file system, cluster file system, or raw devices.

Configuration Requirements

Use the requirements in this section to plan the installation and configuration of Oracle Solaris Cluster HA for Informix. These requirements apply only to Oracle Solaris Cluster HA for Informix. You must meet these requirements before you proceed with your Oracle Solaris Cluster HA for Informix installation and configuration.

Caution – Your data service configuration might not be supported if you do not adhere to these requirements.

For requirements that apply to all data services, see "Configuration Guidelines for Solaris Cluster Data Services" in *Solaris Cluster Data Services Planning and Administration Guide for Solaris OS*.

Determine the Solaris Zone for Informix Dynamic Server to Use

Solaris zones provides a means of creating virtualized operating system environments within an instance of the Solaris 10 OS. Solaris zones enable one or more applications to run in isolation from other activity on your system. For complete information about installing and configuring a Solaris Container, refer to *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*.

You must determine which Solaris zone Informix Dynamic Server will use. Informix Dynamic Server can be deployed in the global zone or non-global zone.

Preparing the Nodes and Disks

Use the files source for the `passwd`, `group`, and `project` database entries in the `/etc/nsswitch.conf` file and do not use external sources such as NIS, NIS+, or LDAP for the `passwd`, `group`, and `project` database entries.

Oracle Solaris Cluster HA for Informix uses the `su` command to start, stop, and probe the Informix application and might be impacted if an external name service is unavailable.

On each node or zone that can master the Informix resource, include the following entries in the `/etc/nsswitch.conf` file.

```
passwd:  files
group:   files
project: files
```

However, you might still use an external name service such as NIS, NIS+, or LDAP. Remember, though, that if the external name service is unavailable, Oracle Solaris Cluster HA for Informix might fail.

Installing and Configuring Informix

This section contains the procedures that you need to install and configure the Informix application.

▼ How to Install Informix

Use this procedure to install the Informix application.

1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.

2 Create a cluster file system or highly available local file system for the Informix files.

Refer to the *Oracle Solaris Cluster Software Installation Guide* for information about creating a cluster file system and the *Oracle Solaris Cluster Data Services Planning and Administration Guide* for information about creating a highly available local file system.

3 Create an Informix user ID and group ID on all the nodes or zones of the cluster that can master the Informix Dynamic Server.

4   Install Informix binaries on the cluster node or zone where the file system created in Step 2 is mounted.

Note – For more information about how to install Informix, see (`http://publib.boulder.ibm.com/infocenter/idshelp/v111/index.jsp`).

Note – Before you install Informix on a cluster file system, use the `scstat`(1M) command to verify that the Solaris Cluster software is fully operational.

Enabling the Informix Instances to Run in a Cluster

This section contains the procedure that you need to verify the Informix installation.

▼ How to Enable the Informix Instances to Run in a Cluster

Use this procedure to verify the Informix instance. Perform the following steps on a node or zone on which the Informix instance can run.

1   Log in to the cluster node or zone on which you installed the Informix software.

2   Shut down the Informix instance.

For more information about how to start or stop Informix, see `http://publib.boulder.ibm.com/infocenter/idshelp/v111/index.jsp`

3   Mount the highly available local file system if used.

4   Create the failover resource group to hold the network and application resources.

```
# clresourcegroup create ids-rg
```

Note – Use the `-n` option to the `clresourcegroup` command to select the set of nodes or zones on which the Informix application can run.

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

---

5   Verify that you have added to your name service database all the network resources that you use.

6   Add a network resource (logical hostname) to the failover resource group.

```
# clreslogicalhostname create -g ids-rg [-h lhost[,...]] [-d] ids-lh
```

7   Edit the `$INFORMIXDIR/etc/sqlhosts.$INFORMIXSERVER` file and update the hostname to reflect the logical host.

8   Add the `HAStoragePlus` resource to the resource group for the failover application.

```
# clresource register SUNW.HAStoragePlus
# clresource create -g ids-rg \
-t SUNW.HAStoragePlus \
-p filesystemmountpoints=mountpoint,...  ids-has
```

The resource is created in the enabled state.

Tip – This HAStoragePlus resource is used for the $INFORMIXSERVER file.

For more information on how to set up an HAStoragePlus resource, see "Enabling Highly Available Local File Systems" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

9   Enable the resource group.

Run the clresourcegroup(1CL) command to move the resource group into a managed state and bring the resource group online.

```
# clresourcegroup online -M ids-rg
```

10   Start the Informix database on a node where the resource group is online.

Ensure that the Informix database starts smoothly by connecting to the database by using the logical hostname.

11   Shut down the Informix database.

12   As user `root`, switch the Informix database resource group to another cluster member.

```
# clresourcegroup switch -n node[:zone]  ids-rg
```

13   Start the Informix database.

Ensure that the Informix database starts smoothly by connecting to the database by using the logical hostname.

14   Ensure that you switch the Informix resource group to all the nodes in the resource group's nodelist and also ensure that you can start Informix on all the nodes of the cluster.

Installing the HA for Informix Packages

If you did not install the HA for Informix packages during your initial Oracle Solaris Cluster installation, perform this procedure to install the packages. To install the packages, use the `installer` program.

Note – You need to install the HA for Informix packages in the global cluster and not in the zone cluster.

▼ How to Install the HA for Informix Packages

Perform this procedure on each cluster node where you are installing the HA for Informix packages.

You can run the `installer` program with a command-line interface (CLI) or with a graphical user interface (GUI). The content and sequence of instructions in the CLI and the GUI are similar.

Note – Even if you plan to configure this data service to run in non-global zones, install the packages for this data service in the global zone. The packages are propagated to any existing non-global zones and to any non-global zones that are created after you install the packages.

Before You Begin   Ensure that you have the Oracle Solaris Cluster installation media.

If you intend to run the `installer` program with a GUI, ensure that your `DISPLAY` environment variable is set.

1   On the cluster node where you are installing the data service packages, become superuser.

2   Load the Oracle Solaris Cluster installation media into the DVD-ROM drive.

If the Volume Management daemon `vold(1M)` is running and configured to manage DVD-ROM devices, the daemon automatically mounts the DVD-ROM on the `cdrom` directory.

3   Change to the installation wizard directory of the DVD-ROM.

■ **If you are installing the data service packages on the SPARC platform, type the following command:**

   `# cd /cdrom/cdrom0/Solaris_sparc`

■ **If you are installing the data service packages on the x86 platform, type the following command:**

   `# cd /cdrom/cdrom0/Solaris_x86`

4   Start the installation wizard.

```
# ./installer
```

5   When you are prompted, accept the license agreement.

6   From the list of Oracle Solaris Cluster agents under Availability Services, select the data service for Informix.

7   If you require support for languages other than English, select the option to install multilingual packages.

English language support is always installed.

8   When prompted whether to configure the data service now or later, choose Configure Later.

Choose Configure Later to perform the configuration after the installation.

9   Follow the instructions on the screen to install the data service packages on the node.

The installation wizard displays the status of the installation. When the installation is complete, the wizard displays an installation summary and the installation logs.

10   (GUI only) If you do not want to register the product and receive product updates, deselect the Product Registration option.

The Product Registration option is not available with the CLI. If you are running the installation wizard with the CLI, omit this step.

11   Exit the installation wizard.

12   Unload the installation media from the DVD-ROM drive.

a.   To ensure that the DVD-ROM is not being used, change to a directory that does *not* reside on the DVD-ROM.

b.   Eject the DVD-ROM.

```
# eject cdrom
```

Next Steps   Go to "Registering and Configuring Oracle Solaris Cluster HA for Informix" on page 19.

Registering and Configuring Oracle Solaris Cluster HA for Informix

This section contains the procedures you need to configure Oracle Solaris Cluster HA for Informix.

▼ How to Register and Configure Oracle Solaris Cluster HA for Informix

Use this procedure to configure Oracle Solaris Cluster HA for Informix.

1 **Become superuser on one of the nodes or zones in the cluster that hosts the Informix database resource group.**

2 **Edit the `/opt/SUNWscids/util/ids_config` file and type values as required by following the comments within the file.**

```
RS=ids-rs
RG=ids-rg
LH_RS=ids-lh
HAS_RS=ids-has
USERID=informix
ONCONFIG=onconfig.demo_on
INFORMIXDIR=/opt/IBM/informix
INFORMIXSERVER=demo_on
INFORMIXSQLHOSTS=/opt/IBM/informix/etc/sqlhosts.demo_on
```

3 **Register the resource type for the failover application.**

```
# clresourcetype register SUNW.gds.
```

4 **Register the Informix resource.**

```
# cd /opt/SUNWscids/util
# ./ids_register
```

Verifying the Oracle Solaris Cluster HA for Informix Installation and Configuration

This section contains the procedure you need to verify that you installed and configured your data service correctly.

▼ How to Verify Installation and Configuration of Oracle Solaris Cluster HA for Informix

Use this procedure to verify the installation and configuration of Oracle Solaris Cluster HA for Informix.

1   Log in to the node or zone that currently hosts the resource group that contains the Informix application resource.

2   As user `root`, switch the Informix resource group to another cluster member.

```
# clresourcegroup switch -n node[:zone] ids-rg
```

3   Verify that the Informix application starts on the node or zone that you specified in Step 1.

4   Repeat Step 1 through Step 3 until you have tested all the potential nodes or zones on which the Informix application can run.

Understanding the Oracle Solaris Cluster HA for Informix Fault Monitor

This section describes the Oracle Solaris Cluster HA for Informix fault-monitor probing algorithm or functionality.

For conceptual information about fault monitors, see the *Sun Cluster Concepts Guide*.

Probing Algorithm and Functionality

Oracle Solaris Cluster HA for Informix fault monitor is controlled by the extension properties that control the probing frequency. The default values of these properties determine the preset behavior of the fault monitor. The preset behavior should be suitable for most Solaris Cluster installations. Therefore, you should tune the Oracle Solaris Cluster HA for Informix fault monitor only if you need to modify this preset behavior. You can modify this behavior by performing the following actions:

- Setting the interval between fault monitor probes (`Thorough_probe_interval`)
- Setting the timeout for fault monitor probes (`Probe_timeout`)
- Setting the number of times the fault monitor attempts to restart the resource (`Retry_count`)

The Oracle Solaris Cluster HA for Informix fault monitor checks the database within an infinite loop. During each cycle the fault monitor checks the database state and reports either a failure or success.

If the fault monitor is successful, it returns to its infinite loop and continues the next cycle of probing and sleeping.

If the fault monitor reports a failure, a request is made to the cluster to restart the resource. If the fault monitor reports another failure, another request is made to the cluster to restart the resource. This behavior continues whenever the fault monitor reports a failure.

If successive restarts exceed the Retry_count within the Thorough_probe_interval, a request is made to fail over the resource group onto a different node or zone.

Operations of the Informix Probe

The Informix probe checks the database state by using the onstat command.

The onstat - command produces a status line for the database server. The status line is accurate at the instant that the command executes and identifies the server and version but more importantly, the status of the database server.

The status modes that are considered to be normal operational modes are On-Line and Read-only. Whenever the onstat command reports these status modes, the Informix probe considers that the database server is operating in an acceptable mode.

The status modes that are considered to be maintenance and startup modes are Quiescent, Administration, Single User, Initialization, Fast Recovery, and Recovery. While the database server is starting, the probe considers these modes acceptable. However, after the database has started successfully, these modes are not considered acceptable and if one of these modes is encountered, the probe requests a restart of the resource.

The probe also requests a resource to restart if any unknown status modes are reported by the onstat command.

Debugging Oracle Solaris Cluster HA for Informix

▼ How to Turn on Debugging for Oracle Solaris Cluster HA for Informix

Oracle Solaris Cluster HA for Informix can be used by multiple Informix instances. Debugging can be turned on for all Informix instances or a particular Informix instance.

The /opt/SUNWscids/etc/config file enables you to turn on debugging for all Informix instances or for a specific Informix instance on a particular node or zone within the cluster. If you require debugging to be turned on for Oracle Solaris Cluster HA for Informix across the whole cluster, repeat this step on all nodes within the cluster.

1 Edit the `/etc/syslog.conf` file and change `daemon.notice` to `daemon.debug`.

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

Change the `daemon.notice` file to `daemon.debug` and restart the `syslogd` command. Note that the following output, from `grep daemon /etc/syslog.conf`, 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
```

Restart the syslog daemon.

- **If you are running Solaris 10, type:**

  ```
  # svcadm disable system-log
  # svcadm enable system-log
  ```

2 Edit the `/opt/SUNWscids/etc/config` file.

Perform this step for each component that requires debugging output, on each node of Solaris Cluster as required.

Edit the /opt/SUNWscids/etc/config file and change DEBUG= to DEBUG=ALL or DEBUG=*sun-cluster-resource*.

```
# cat /opt/SUNWscids/etc/config
#
# Copyright 2006 Sun Microsystems, Inc.  All rights reserved.
# Use is subject to license terms.
#
# ident "@(#)config     1.1      06/03/06 SMI"
#
# Usage:
#        DEBUG=<RESOURCE_NAME> or ALL
#
DEBUG=ALL
```

Note – To turn off debug, reverse the previous steps.

Index