Oracle® Solaris Cluster Data Service for Sun Java System Application Server EE (HADB) Guide

Part No: 821–1542–10
September 2010, Revision A

ORACLE

Copyright © 2000, 2010, Oracle and/or its affiliates. All rights reserved.

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

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

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

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

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

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

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd.

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

110714@25097

Contents

Preface

Oracle Solaris Cluster Data Service for Sun Java System Application Server EE (HADB) Guide explains how to install and configure Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB).

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 Sun 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.

◆ ◆ ◆ CHAPTER 1

Oracle Solaris Cluster HA for Sun Java System Application Server EE (Supporting HADB Versions as of 4.4)

This document describes the procedures to install and configure the data service called Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB).

This document contains two chapters.

- This chapter contains information for the data service when it makes highly available those versions of the Sun Java System Application Server EE (HADB) application *as of* version 4.4 that are supported by the Oracle Solaris Cluster product and by this data service.
- Chapter 2, "SPARC: Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB Versions Before 4.4)," contains information for the data service when it makes highly available those versions of the Sun Java System Application Server EE (HADB) application *before* version 4.4 that are supported by the Oracle Solaris Cluster product and by this data service.

This chapter contains the following sections.

- "Tuning the Fault Monitor for Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)" on page 26

Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) Overview

This section describes how the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) data service enables the Management Agent component of Sun Java System Application Server EE (HADB) to be highly available in an Oracle Solaris Cluster system. The data service must be configured to be mastered by multiple nodes.

See Chapter 1, "Planning for Oracle Solaris Cluster Data Services," in *Oracle Solaris Cluster Data Services Planning and Administration Guide* for general information about data services. The data service must be configured

The Sun Java System Application Server EE (HADB) software is packaged with the Sun Java System Application Server Enterprise Edition installation. However, you can run Sun Java System Application Server EE (HADB) and Sun Java System Application Server on separate clusters if appropriate. This document describes how to install and configure the data service that enables you to use Sun Java System Application Server EE (HADB) in a cluster. First, enable Sun Java System Application Server EE (HADB) in your cluster to provide session and Enterprise Java Bean (EJB) persistence. See *Sun Java System Application Server Enterprise Edition 8.2 High Availability Administration Guide* for information about Sun Java System Application Server EE (HADB).

Implementation of the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) does not assume the existence of applications on which your architecture depends, such as databases and web servers. However, such applications can be configured to be highly available and might run on a different cluster.

HADB Management Architecture

The following figure illustrates the architecture of a database with four active nodes in a management domain.

FIGURE 1–1 Management System Architecture

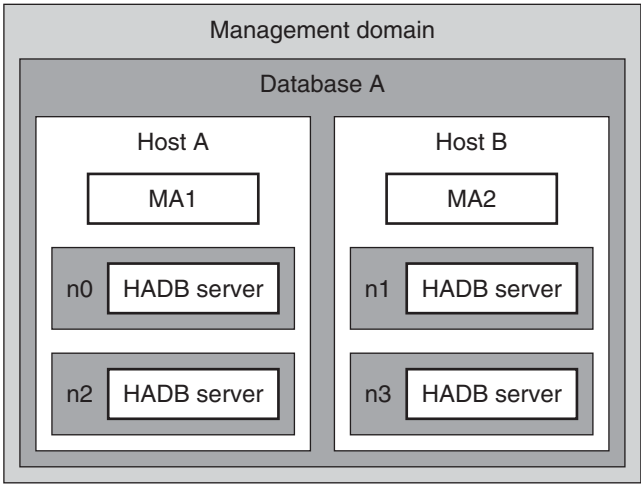

The figure illustrates a management domain that consists of a database and four HADB node processes in the database. These node processes are indicated by n0, n1, n2, and n3. Host A and Host B represent the cluster nodes and form the Data Redundancy Units, or DRUs (Host A:DRU0 and Host B:DRU1) for HADB.

Management agent MA1 manages n0 and n2, and management agent MA2 manages n1 and n3. The management agent is responsible for monitoring the health of the node and restarting the node when necessary. The management agent also handles the management operations requests from the clients.

If the management agent fails on the cluster node, the agent must be restarted on the same node so that the agent can continue to monitor the HADB nodes and to provide management operations. Therefore, the management agent must be made highly available.

Highly Available HADB Management Agent

The Management Agent component of Sun Java System Application Server EE (HADB) is made highly available by creating a data service resource in a resource group that is configured to be mastered by multiple nodes at the same time. In this resource group, the maximum number of primary nodes and the desired number of primary nodes must be the same. This configuration is described in detail in "Registering and Configuring the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)" on page 22.

The following figure illustrates a two-node configuration before any failure occurs on a node.

FIGURE 1–2 Management Agents Mastered on Multiple Nodes—Before Failure

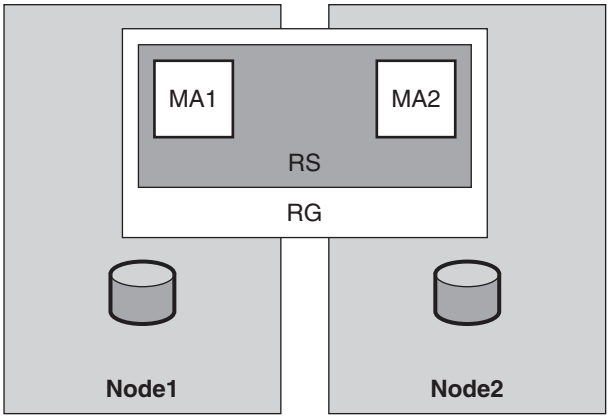

The figure illustrates two management agents, indicated by MA1 and MA2, on different cluster nodes. The application is installed on the local file system on each node. With the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) data service, the two management agents are configured in one resource, indicated by RS, and this resource is configured in one resource group, RG. The resource group is configured to be mastered on multiple cluster nodes at one time, in this case, Node1 and Node2.

The data service manages starting the management agents on all the nodes on which they are configured. The management agents in turn start the Node Supervisor (NSUP) process, which then starts the HADB node processes. If a management agent fails, the data service restarts it, based on user-defined parameters. If the HADB node processes fail, the management agent restarts them. If any databases have stopped, the user must restart these databases by issuing commands on the command line.

If the user-defined parameters determine that the data service does not restart the management agent after a failure, the processes remain down on that node. However, the management agent processes on the other managed nodes continue to execute. No failover occurs.

The following figure illustrates the two-node configuration after the failure of the first node. All the processes on the first node are down, and the second node continues to execute its processes.

FIGURE 1–3 Management Agents Mastered on Multiple Nodes—After Failure

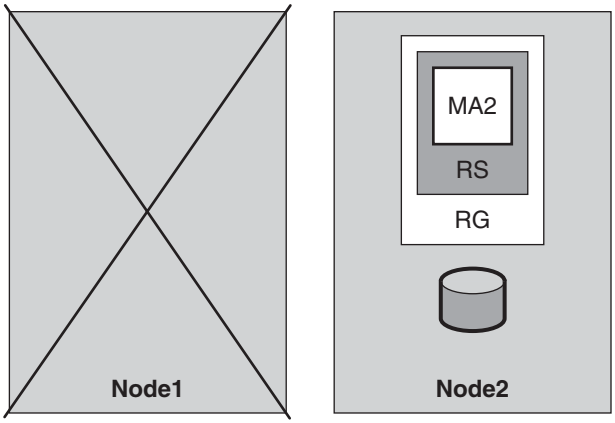

Overview of the Installation and Configuration Process for Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)

The following table summarizes the tasks for installing and configuring Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) and provides cross-references to detailed instructions for performing these tasks. Perform the tasks in the order that they are listed in the table.

TABLE 1–1 Tasks for Installing and Configuring Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)

Task	For Instructions
Plan the installation and configuration of Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)	"Planning the Installation and Configuration of Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)" on page 18
Install and configure the Sun Java System Application Server EE (HADB) software	"Installing and Configuring the Sun Java System Application Server EE (HADB) Software" on page 19
Create a Sun Java System Application Server EE (HADB) management domain	"Creating a Sun Java System Application Server EE (HADB) Management Domain" on page 20
Install the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) package	"How to Install the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) Packages" on page 21

TABLE 1–1 Tasks for Installing and Configuring Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) *(Continued)*

Task	For Instructions
Register and configure Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) as a data service mastered on multiple nodes at one time	"Registering and Configuring the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)" on page 22
Verify the installation and configuration of Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)	"Verifying the Installation and Configuration of Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)" on page 25
Tune the fault monitor for Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)	"Tuning the Fault Monitor for Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)" on page 26

Note – If you run multiple data services in your Oracle Solaris Cluster configuration, you can set up the data services in any order, with the following exception. If Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) depends on the Oracle Solaris Cluster HA for DNS, you must set up DNS first. For details, see *Oracle Solaris Cluster Data Service for Domain Name Service (DNS) Guide*. DNS software is included in the Solaris software. If the cluster is to obtain the DNS service from another server, configure the cluster to be a DNS client first.

Planning the Installation and Configuration of Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)

This section contains the information that you need to plan your installation and configuration of Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB).

Configuration Restrictions and Requirements

Store static files and data on the local file system of each cluster node. The Sun Java System Application Server EE (HADB) software is installed when you install the Sun Java System Application Server Enterprise Edition software. See *Oracle Solaris Cluster Data Service for Sun Java System Application Server Guide* documentation for instructions. When the database is created, the configuration and data files are created by default on the local file system of each cluster node. For details, see the Sun Java System Application Server documentation.

If you are using Solaris 5.8, you must perform the following procedure.

1. Edit the file `/etc/init.d/ma-initd`.

2. Add the following entry at the beginning of the file.

```
_JAVA_OPTIONS="-Djava.net.preferIPv4Stack=true"
export _JAVA_OPTIONS
```

3. Save the file.

Configuration Planning Considerations

For information to help you plan the installation and configuration of Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB), see "Considerations for Installing and Configuring a Data Service" in *Oracle Solaris Cluster Data Services Planning and Administration Guide*. See also the worksheets in Appendix C, "Data Service Configuration Worksheets and Examples," in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

If you choose to use your HADB data service with another highly available application, resource dependencies might exist. See Appendix A, "Standard Properties," in *Oracle Solaris Cluster Data Services Planning and Administration Guide* for a description of the `Resource_dependencies` property.

Installing and Configuring the Sun Java System Application Server EE (HADB) Software

The Sun Java System Application Server EE (HADB) software is a Java 2 Enterprise Edition (J2EE™) 1.3 compliant relational database. The Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) data service is designed to meet the needs of enterprise customers and run under the control of Solaris Cluster software. The Sun Java System Application Server provides a transactional session-state persistence infrastructure that is highly available and highly scalable. Application Server uses the HADB to store session information. The HADB management client is the command-line interface for the HADB. A full set of utilities are available for performing HADB configuration, runtime management, and monitoring.

Instructions for using these utilities are contained in the Sun Java System Application Server documentation, the `hadbm` man pages, and the `asadmin` command session-persistence man pages. For information about installing and configuring HADB with Sun Java System Application Server, see the Sun Java System Application Server documentation. For information about configuring Solaris Cluster HA for Sun Java System Application Server, see *Oracle Solaris Cluster Data Service for Sun Java System Application Server Guide*.

Note – If you are using the Solaris 10 OS, install and configure this data service to run only in the global zone. At publication of this document, this data service is not supported in non-global zones. For updated information about supported configurations of this data service, contact your Oracle service representative.

Creating a Sun Java System Application Server EE (HADB) Management Domain

Before putting the HADB management agents under control of the data service, you must create a Sun Java System Application Server EE (HADB) management domain.

▼ Creating a Sun Java System Application Server EE (HADB) Management Domain

Use the example in the following procedure to create, start, and verify a management domain.

For details, see the Sun Java System Application Server documentation.

1 Create the management domain.

Use the `hadbm` command with the `createdomain` option. This command automatically starts the domain.

2 Verify that the domain is created and is running.

Use the `hadbm` command with the `listdomain` option.

Installing the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) Packages

If you did not install the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) 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 Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) packages in the global cluster and not in the zone cluster.

▼ How to Install the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) Packages

Perform this procedure on each cluster node where you are installing the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) 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.

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 Sun Java System Application Server EE (HADB).**

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 See "Registering and Configuring Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)" on page 34 to register Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) and to configure the cluster for the data service.

Registering and Configuring the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)

This procedure describes how to use the Oracle Solaris Cluster maintenance commands to register and configure Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) as a resource mastered on multiple nodes at one time.

Note – See "Tools for Data Service Resource Administration" in *Oracle Solaris Cluster Data Services Planning and Administration Guide* for details about additional options that enable you to register and configure the data service.

To perform this procedure, you need the following information about your configuration.

- The name of the resource type for Oracle Solaris Cluster HA for Sun Java System Application Server EE (supporting HADB versions as of 4.4), which is `SUNW.hadb_ma`.

- The names of the cluster nodes that can master the data service.

Setting Extension Properties for Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)

The sections that follow contain instructions for registering and configuring Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) resources. For information about the extension properties, see Appendix A, "Properties for Oracle Solaris Cluster HA for Sun Java System Application Server EE (Supporting HADB Versions as of 4.4)." The Tunable entry indicates when you can update a property.

See Appendix A, "Standard Properties," in *Oracle Solaris Cluster Data Services Planning and Administration Guide* for details about all of the Oracle Solaris Cluster properties.

To set an extension property of a resource, include the following option in the `clresource`(1CL) command that creates or modifies the resource:

`-p` *property*=*value*

`-p` *property*
: Identifies the extension property that you are setting.

value
: Specifies the value to which you are setting the extension property.

You can also use the procedures in Chapter 2, "Administering Data Service Resources," in *Oracle Solaris Cluster Data Services Planning and Administration Guide* to configure resources after the resources are created.

▼ How to Register and Configure Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)

Perform the following steps to complete your configuration.

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

2. **Register the resource type for Sun Java System Application Server EE (HADB).**

```
# clresourcetype register SUNW.hadb_ma
```

3 Create the resource group for the Sun Java System Application Server EE (HADB) software.

This resource group is configured to be mastered by multiple nodes at the same time.

```
# clresourcegroup create -n nodelist \
-p Maximum_primaries=nodes_in_rg \
-p Desired_primaries=nodes_in_rg resource_group
```

`-n` *nodelist*
Specifies a comma-separated subset of cluster nodes to run the Sun Java System Application Server EE (HADB) software. If this option is omitted, all cluster nodes run Sun Java System Application Server EE (HADB). Use the command to find the node list names.

`-p Maximum_primaries=`*nodes_in_rg*
Specifies the maximum number of nodes on which the resource can start. You must specify the same number as the value of the `Desired_primaries` property.

`-p Desired_primaries=`*nodes_in_rg*
Specifies the desired number of nodes on which the resource can start. You must specify the same number as the value of the `Maximum_primaries` property.

resource_group
Specifies the name of the resource group that is to be created.

4 Create a resource for Sun Java System Application Server EE (HADB).

See Appendix A, "Properties for Oracle Solaris Cluster HA for Sun Java System Application Server EE (Supporting HADB Versions as of 4.4)," for a detailed description of the extension properties for this resource type.

```
# clresource create -g resource_group  -t SUNW.hadb_ma \
-p HADBM_PASSWORDFILE=path_to_pw_file resource
```

`-g` *resource_group*
Specifies the name of the resource group to which the resource is to be added.

`-t SUNW.hadb_ma`
Specifies the predefined resource type name.

`-p HADBM_PASSWORDFILE=`*path_to_pw_file*
Specifies the full path to the file containing the HADBM administrative password.

For details about the contents of the password file, see the Sun Java System Application Server documentation.

resource
Specifies the name of the resource that you are creating.

The resource is created in the enabled state.

5 Bring the resource group online.

```
# clresourcegroup online resource_group
```

6 **Verify that the resource group and HADB resource are online.**

```
# cluster status
# ps -ef
```

7 **To verify that you have correctly installed and configured Solaris Cluster HA for HADB, run the following command.**

```
# hadbm status database_name --nodes
```

The output should indicate that the database that you specified is running.

Example 1–1 Creating a `SUNW.hadb_ma` Resource

This example shows the creation of a `SUNW.hadb_ma` resource.

In this example the default values for the extension properties are used. The resource has the following characteristics:

- The resource is named `hadb-rs`.
- The resource is a member of a resource group named `hadb-rg`.
- The resource is an instance of the `SUNW.hadb_ma` resource type. The registration of the resource type is not shown in this example.

```
clresource create -g hadb-rg -t SUNW.hadb_ma \
-p HADBM_PASSWORDFILE=/global/disk1/hadbm-pwfile hadb-rs
```

Verifying the Installation and Configuration of Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)

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

▼ How to Verify the Installation and Configuration of Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)

1 **Ensure that the Sun Java System Application Server EE (HADB) software is started under the control of Solaris Cluster software.**

```
# clresourcegroup online -M resource_group
```

2 **Verify that the resource group and Sun Java System Application Server EE (HADB) resource are online.**

```
# clresourcegroup status
# ps -ef
```

3 **Verify that you have correctly installed and configured Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB).**

```
# hadbm status database_name --nodes
```

The output should indicate that the database that you specified is running.

Tuning the Fault Monitor for Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)

The fault monitor for Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) is contained in the resource that represents Sun Java System Application Server EE (HADB). You create this resource when you register and configure Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB). For more information, see "Registering and Configuring the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)" on page 22.

System properties and extension properties of this resource control the behavior of the fault monitor. The default values of these properties determine the preset behavior of the fault monitor. The preset behavior should be suitable for most Oracle Solaris Cluster installations. Therefore, you should tune the fault monitor *only* if you need to modify this preset behavior.

For more information, see the following sections.

- "Tuning Fault Monitors for Oracle Solaris Cluster Data Services" in *Oracle Solaris Cluster Data Services Planning and Administration Guide*
- "Changing Resource Type, Resource Group, and Resource Properties" in *Oracle Solaris Cluster Data Services Planning and Administration Guide*
- Appendix A, "Standard Properties," in *Oracle Solaris Cluster Data Services Planning and Administration Guide*
- Appendix A, "Properties for Oracle Solaris Cluster HA for Sun Java System Application Server EE (Supporting HADB Versions as of 4.4)," in this guide for a detailed description of the extension properties for the `SUNW.hadb_ma` resource type

◆ ◆ ◆ CHAPTER 2

SPARC: Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB Versions Before 4.4)

This document describes the procedures to install and configure the data service called Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB).

This document contains two chapters.

- This chapter contains information for the data service when it makes highly available those versions of the Sun Java System Application Server EE (HADB) application *before* version 4.4 that are supported by the Oracle Solaris Cluster product and by this data service.
- Chapter 1, "Oracle Solaris Cluster HA for Sun Java System Application Server EE (Supporting HADB Versions as of 4.4)," contains information for the data service when it makes highly available those versions of the Sun Java System Application Server EE (HADB) application *as of* version 4.4 that are supported by the Oracle Solaris Cluster product and by this data service.

This chapter contains the following sections.

- "Tuning the Fault Monitor for Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)" on page 39

Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) Overview

This section describes how the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) data service enables the Sun Java System Application Server EE (HADB) software for high availability.

The Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) is a data service with appropriate extension properties to configure a database that is mastered on multiple nodes at one time.

See Chapter 1, "Planning for Oracle Solaris Cluster Data Services," in *Oracle Solaris Cluster Data Services Planning and Administration Guide* for general information about data services.

The Sun Java System Application Server EE (HADB) software is packaged with the Sun Java System Application Server Enterprise Edition installation. However, you can run Sun Java System Application Server EE (HADB) and Sun Java System Application Server on separate clusters if appropriate. This document describes how to install and configure the data service that enables you to use Sun Java System Application Server EE (HADB) in a cluster. First, enable Sun Java System Application Server EE (HADB) in your cluster to provide session and Enterprise Java Bean (EJB) persistence. See *Sun Java System Application Server Enterprise Edition 8.2 High Availability Administration Guide* for information about the Sun Java System Application Server EE (HADB) software.

Implementation of the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) does not assume the existence of applications on which your architecture depends, such as databases and web servers. However, such applications can be configured to be highly available and might run on a different cluster.

Overview of the Installation and Configuration Process for Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)

The following table summarizes the tasks for installing and configuring Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) and provides cross-references to detailed instructions for performing these tasks. Perform the tasks in the order that they are listed in the table.

TABLE 2–1 Tasks for Installing and Configuring Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)

Task	For Instructions
Plan the installation and configuration of Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)	"Planning the Installation and Configuration of Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)" on page 30
Install and configure the Sun Java System Application Server EE (HADB) software	"Installing and Configuring the Sun Java System Application Server EE (HADB) Software" on page 32
Create a Sun Java System Application Server EE (HADB) database	"How to Create a Sun Java System Application Server EE (HADB) Database" on page 33
Install the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) package	"Installing the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) Packages" on page 33
Register and configure Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) as a data service mastered on multiple nodes at one time	"How to Register and Configure Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)" on page 35
Verify the installation and configuration of Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)	"Verifying the Installation and Configuration of Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)" on page 37
Maintain the HADB database	"Maintaining the HADB Database" on page 38
Tune the fault monitor for Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)	"Tuning the Fault Monitor for Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)" on page 39

Note – If you run multiple data services in your Oracle Solaris Cluster configuration, you can set up the data services in any order, with the following exception. If the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) data service depends on the Oracle Solaris Cluster HA for DNS data service, you must set up DNS first. For details, see *Oracle Solaris Cluster Data Service for Domain Name Service (DNS) Guide*. The DNS software is included in the Solaris software. If the cluster is to obtain the DNS service from another server, you should first configure the cluster to be a DNS client.

Planning the Installation and Configuration of Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)

This section contains the information that you need to plan your installation and configuration of Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB).

Store static files and data on the local file system of each cluster node. The Sun Java System Application Server EE (HADB) software is installed when you install the Sun Java System Application Server Enterprise Edition software. See *Oracle Solaris Cluster Data Service for Sun Java System Application Server Guide* documentation for instructions. When the database is created, the configuration and data files are created by default on the local file system of each cluster node. For details, see the Sun Java System Application Server documentation.

Configuration Overview

Use the multiple masters configuration in this section to plan the installation and configuration of the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) data service. This data service might support additional configurations. However, you must contact your Enterprise Services representative for information about additional configurations.

HADB configurations are defined by data nodes, or a collection of processes. Each node is a dedicated area of main memory with one or more secondary storage devices. These storage devices are not shared storage. Each HADB data node must have exclusive access to an area of main memory and several areas of disk space. HADB data nodes are active or spare.

The recommended minimum requirement for a cluster running the Sun Java System Application Server EE (HADB) and the Sun Java System Application Server is four active data nodes plus two spare nodes. See "Creating a Sun Java System Application Server EE (HADB) Database" on page 32 for an example of how to create a Sun Java System Application Server EE (HADB) with four active nodes and two spare nodes by using the `hadbm` comand-line utility. For high availability, you configure the HADB data redundancy unit (DRU) to use the Solaris Cluster interconnect. See the Sun Java System Application Server documentation for detailed information about DRUs. The HADB JDBC driver in the client application handles highly available access to the database.

The following figure illustrates the recommended minimum configuration for the Sun Java System Application Server EE (HADB) and the Sun Java System Application Server.

FIGURE 2–1 Recommended Minimum Configuration

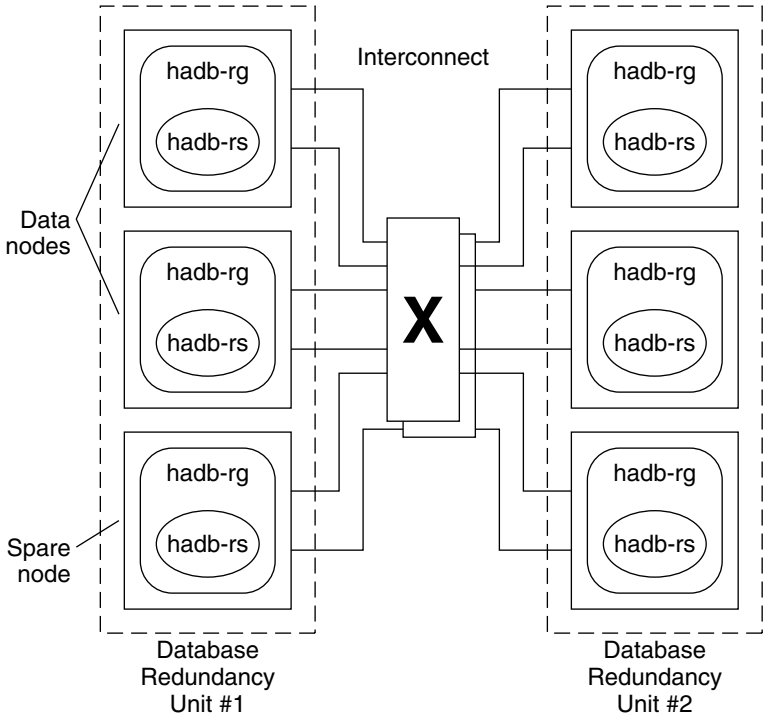

Configuration Planning Questions

Use the questions in this section to plan the installation and configuration of the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) data service. See "Considerations for Installing and Configuring a Data Service" in *Oracle Solaris Cluster Data Services Planning and Administration Guide* for information that might apply to these questions. See also the worksheets in Appendix C, "Data Service Configuration Worksheets and Examples," in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

If you choose to use your HADB data service with another highly available application, resource dependencies might exist. See Appendix A, "Standard Properties," in *Oracle Solaris Cluster Data Services Planning and Administration Guide* for a description of the `Resource_dependencies` property.

- Will all history files, data and log devices, and the database configuration files reside on local file systems?
- What are the Oracle Solaris Cluster private interconnect hostnames that you will use?

Installing and Configuring the Sun Java System Application Server EE (HADB) Software

The Sun Java System Application Server EE (HADB) software is a Java 2 Enterprise Edition (J2EE™) 1.3 compliant relational database. The Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) data service is designed to meet the needs of enterprise customers and run under the control of Solaris Cluster software. The Sun Java System Application Server provides a transactional session-state persistence infrastructure that is highly available and highly scalable. Application Server uses the HADB to store session information. The HADB management client is the command-line interface (CLI) for the HADB. A full set of utilities are available for performing HADB configuration, runtime management, and monitoring.

Instructions for using these utilities are contained in the Sun Java System Application Server documentation, the `hadbm` man pages, and the `asadmin` command session-persistence man pages. The Sun Java System Application Server EE (HADB) software is co-packaged with the Sun Java System Application Server Enterprise Edition. For information about installing and configuring HADB with Sun Java System Application Server, see *Sun Java System Application Server Enterprise Edition 8.2 High Availability Administration Guide*. For information about configuring Oracle Solaris Cluster HA for Sun Java System Application Server, see *Oracle Solaris Cluster Data Service for Sun Java System Application Server Guide*.

Creating a Sun Java System Application Server EE (HADB) Database

This section contains the procedure to configure and create the initial HADB database in a Oracle Solaris Cluster environment. Consider the following restrictions before creating your database.

- You must specify an even number of Oracle Solaris Cluster hosts by using the `--hosts` option during database creation.
- You must specify the Oracle Solaris Cluster hosts by using the Solaris Cluster private interconnect hostnames. An example of a private interconnect hostname is `clusternode1-priv`.
- Sun Java System Application Server EE (HADB) mirror nodes must be located on different Oracle Solaris Cluster nodes.
- Do not use the `--inetd` option.
- All history files, data and log devices, and the database configuration files must reside on local file systems.
- If more than one Solaris Cluster node is stopped, the entire database is shut down. You must either stop only one node at a time or stop all nodes. You can stop all nodes with the `clresourcegroup offline` command.

- The HADB resource group must not be used with HA Storage Plus.

How to Create a Sun Java System Application Server EE (HADB) Database

Use the example in the following procedure to create, start, and verify the database.

1. **Create the database. This command automatically starts the database.**

```
# hadbm create \
-H clusternode1-priv,clusternode2-priv,clusternode3-priv, \
clusternode4-priv,clusternode5-priv,clusternode6-priv \
--devicesize=2048 \
-a 4 --set ManagementProtocol=rsh --dbpassword=secret12 \
-s 2 hadb
```

For details, see the Sun Java System Application Server documentation.

Note – You must specify the hosts by using the Oracle Solaris Cluster private interconnect hostnames. To find these hostnames, run the `clnode show` command from the primary Solaris Cluster node.

Note – If you are using the recommended SSH setup, you do not need to specify the `ManagementProtocol` property.

---

2. **Verify that the database is running.**

```
# hadbm status hadb --nodes
```

3. **Stop the database.**

```
# hadbm stop hadb
```

4. **Create the session store and the JDBC connection pool.**

For details, see the Sun Java System Application Server documentation.

Installing the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) Packages

See "Installing the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) Packages" on page 20 to install Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) package.

Registering and Configuring Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)

This procedure describes how to use the Oracle Solaris Cluster maintenance commands to register and configure the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) data service as a resource mastered on multiple nodes at one time.

Note – See "Tools for Data Service Resource Administration" in *Oracle Solaris Cluster Data Services Planning and Administration Guide* for details about additional options that enable you to register and configure the data service.

To perform this procedure, you need the following information about your configuration.

- The name of the resource type for Oracle Solaris Cluster HA for Sun Java System Application Server EE (supporting HADB versions before 4.4), which is `SUNW.hadb`
- The names of the cluster nodes that can master the data service

Setting Extension Properties for Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)

The sections that follow contain instructions for registering and configuring resources for Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB). For information about the extension properties, see Appendix B, "Properties for Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB Versions Before 4.4)." The Tunable entry indicates when you can update a property.

See Appendix A, "Standard Properties," in *Oracle Solaris Cluster Data Services Planning and Administration Guide* for details about all of the Oracle Solaris Cluster properties.

To set an extension property of a resource, include the following option in the `clresource(1CL)` command that creates or modifies the resource:

`-p` *property*=*value*

`-p` *property*
: Identifies the extension property that you are setting.

value
: Specifies the value to which you are setting the extension property.

You can also use the procedures in Chapter 2, "Administering Data Service Resources," in *Oracle Solaris Cluster Data Services Planning and Administration Guide* to configure resources after the resources are created.

▼ How to Register and Configure Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)

Perform the following steps to complete your configuration.

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

2 Register the resource type for the Sun Java System Application Server EE (HADB).

```
# clresourcetype register SUNW.hadb
```

3 Create the resource group for the Sun Java System Application Server EE (HADB).

```
# clresourcegroup create -n nodelist \
-p Maximum_primaries=nodes_in_rg \
-p Desired_primaries=nodes_in_rg resource_group
```

`-n` *nodelist*
Specifies a comma-separated subset of cluster nodes to run Sun Java System Application Server EE (HADB). If this option is omitted, all cluster nodes run Sun Java System Application Server EE (HADB). Use the `clnode list` command to find the node list names.

`-p Maximum_primaries=`*nodes_in_rg*
Specifies the maximum number of nodes on which the resource can start. You must specify the same number as the value of the `Desired_primaries` property.

`-p Desired_primaries=`*nodes_in_rg*
Specifies the desired number of nodes on which the resource can start. You must specify the same number as the value of the `Maximum_primaries` property.

4 Create a resource for Sun Java System Application Server EE (HADB), either with or without auto-recovery.

- If you do not require the auto-recovery feature, execute the following command.

```
# clresource create -g resource_group -t SUNW.hadb \
-p Confdir_list=config_directory_list \
-p HADB_ROOT=install_directory \
-p DB_name=database_name resource
```

`-g` *resource_group*
Specifies that the resource is to be added to the resource group named *resource_group*.

`-t SUNW.hadb`
Specifies the predefined resource type name.

`-p Confdir_list=`*config_directory_list*
Specifies the path to the configuration directory.

`-p HADB_ROOT=`*install_directory*
Specifies the installation directory.

`-p DB_name=`*database_name*
Specifies the name of the HADB database.

The resource is created in the enabled state.

- If you *do* want to use the auto-recovery feature, execute the following command.

```
# clresource create -g resource_group -t SUNW.hadb \
-p Confdir_list=config_directory_list \
-p HADB_ROOT=install_directory \
-p DB_name=database_name \
-p Auto_recovery=TRUE \
-p Auto_recovery_command=command \
-p DB_password_file=password_file resource
```

`-p Auto_recovery=TRUE`
Specifies that you are using the auto-recovery feature.

`-p Auto_recovery_command=`*command*
Specifies the command to execute after the database recovery. This extension property is optional, regardless of the value of the `Auto_recovery` property.

`-p DB_password_file=`*password_file*
Specifies the file from which HADB reads the database password. See the Sun Java System Application Server documentation for the format and content of the password file.

The resource is created in the enabled state.

Note – The auto-recovery command and the database password file must both exist in the local file system on each node.

---

5 Bring the resource group online.

```
# clresourcegroup online resource_group
```

resource_group
Specifies the name of the application resource group that is being enabled.

6 Verify that the resource group and HADB resource are online.

```
# cluster status
# ps -ef
```

7 To verify that you have correctly installed and configured Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB), run the following command.

```
# hadbm status database_name --nodes
```

The output should indicate that the database that you specified is running.

Example 2–1 Creating a SUNW.hadb Resource With Auto-Recovery

This example shows the creation of a SUNW.hadb resource with auto-recovery.

In this example the resource has the following characteristics:

- The resource is named hadb-rs.
- The resource is a member of a resource group named hadb-rg.
- The resource is an instance of the SUNW.hadb resource type. The registration of the resource type is not shown in this example.
- The configuration directory is located at /etc/opt/SUNWhadb/dbdef/hadb.
- The installation directory is located at /opt/SUNWappserver7/SUNWhadb/4.
- The HADB database instance associated with this resource is named hadb.
- Auto-recovery is on.
- The full path to the command to execute after auto-recovery is /usr/local/etc/create-session-store.
- The HADB password file is /usr/local/etc/hadb-password-file.

```
clresource create -g hadb-rg -t SUNW.hadb \
-p Confdir_list=/etc/opt/SUNWhadb/dbdef/hadb  \
-p HADB_ROOT=/opt/SUNWappserver7/SUNWhadb/4 \
-p DB_name=hadb \
-p Auto_recovery=true \
-p Auto_recovery_command=/usr/local/etc/create-session-store \
-p DB_password_file=/usr/local/etc/hadb-password-file hadb-rs
```

Verifying the Installation and Configuration of Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)

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

▼ How to Verify the Installation and Configuration for Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)

1 Ensure that the HADB is started under the control of Oracle Solaris Cluster software.

```
# clresourcegroup online -M resource_group
```

2 **Verify that the resource group and HADB resource are online.**

```
# cluster status
# ps -ef
```

3 **Verify that you have correctly installed and configured Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB).**

```
# hadbm status database_name --nodes
```

The output should indicate that the database that you specified is running.

Maintaining the HADB Database

This section explains how to maintain the HADB database within the Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) data service.

▼ How to Maintain the HADB Database

When you want to run HADB maintenance commands that initiate a rolling restart of the HADB nodes, the HADB resource probe in the fault monitor must be disabled before the maintenance commands are executed and then enabled after the commands and the rolling restart have been completed.

1 **Disable the fault monitor.**

```
# clresource unmonitor resource
```

2 **Run the commands that might initiate a rolling restart.**

The `hadbm` subcommands that might result in a rolling restart include `set`, `restart`, and `addnodes`.

3 **Reenable the fault monitor.**

```
# clresource monitor resource
```

More Information Alternative Method of Maintaining the Database

Alternatively, the HADB resource can be disabled and HADB started outside Oracle Solaris Cluster control while maintenance commands are being performed.

Tuning the Fault Monitor for Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)

The fault monitor for Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB) is contained in the resource that represents Sun Java System Application Server EE (HADB). You create this resource when you register and configure Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB). For more information, see "Registering and Configuring Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB)" on page 34.

System properties and extension properties of this resource control the behavior of the fault monitor. The default values of these properties determine the preset behavior of the fault monitor. The preset behavior should be suitable for most Oracle Solaris Cluster installations. Therefore, you should tune the fault monitor *only* if you need to modify this preset behavior.

For more information, see the following sections.

- "Tuning Fault Monitors for Oracle Solaris Cluster Data Services" in *Oracle Solaris Cluster Data Services Planning and Administration Guide*
- "Changing Resource Type, Resource Group, and Resource Properties" in *Oracle Solaris Cluster Data Services Planning and Administration Guide*
- Appendix A, "Standard Properties," in *Oracle Solaris Cluster Data Services Planning and Administration Guide*

Operations by the Fault Monitor During a Probe

The start method of the HADB resource starts HADB nodes that are configured to run on the local Oracle Solaris Cluster node if they are not running. The method then attempts to start the HADB database. If a failure occurs, the database is started later in the probe.

The fault monitor probe periodically checks the status of the HADB database and the HADB nodes. The probe restarts failed HADB nodes. The probe also starts the HADB database if the HADB resource is not ready to start the database during the start method. For each iteration of this procedure, the probe executes the following steps:

1. The probe retrieves the current status of the HADB database and the HADB nodes by executing the `hadbm status` and `hadbm status --nodes` commands.
2. If the database is not running, the probe checks that the HADB `stopstate` file corresponding to that database exists on the local Oracle Solaris Cluster node. The `hadbm start` command references the `stopstate` file for role assignment of nodes when it starts the database.
3. If the `stopstate` file exists, the HADB resource examines it to determine if the database can be started.

- If the database can be started, the probe starts the database and sets the resource status to `Online`.
- If the database cannot be started, the probe sets the resource status to `Online Degraded`.

4. If the database is running, the probe starts the HADB nodes configured to run on the local Oracle Solaris Cluster node.
5. If the database and the local HADB nodes are running, the probe sets the resource status to `Online` if it was `Online Degraded`.
6. If all the Oracle Solaris Cluster nodes in the HADB resource group have the HADB resource running in the `Online Degraded` state longer than `Stop_timeout` seconds, the HADB resource concludes that the database cannot be started. For a description of the `Stop_timeout` property, see the *method*`_timeout` resource property in Appendix A, "Standard Properties," in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.
7. If the `Auto_recovery` extension property is set to `TRUE`, the HADB resource attempts to recover the database.
8. If recovery of the database is attempted, the probe executes the following steps:
   - Issues the `hadbm clear --fast` command on one of the Oracle Solaris Cluster nodes in the resource group's node list. This command clears the database contents and reinitializes and restarts the database.
   - If the `hadbm clear` command succeeds, the command specified in `Auto_recovery_command` is issued on the same Solaris Cluster node that issued the `hadbm clear` command. The command would normally be a script that contains the `asadmin create-session-store` command. The command can also perform other actions. For example, it might send mail to the Application Server administrator.
   - If both steps succeed, the probe sets the state of the resource to `Online`.

◆ ◆ ◆ APPENDIX A

Properties for Oracle Solaris Cluster HA for Sun Java System Application Server EE (Supporting HADB Versions as of 4.4)

This section describes the extension properties for the resource type `SUNW.hadb_ma`. This resource type represents the Sun Java System Application Server EE (HADB) application (supporting HADB versions as of 4.4) in a Oracle Solaris Cluster configuration.

For details about system-defined properties, see the `r_properties(5)` man page and the `rg_properties(5)` man page.

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

`HADB_MA_CFG`
The full path to the configuration file that is used to start the HADB Management Agent Server.

Data type String

Default `/etc/opt/SUNWhadb/mgt.cfg`

Range Not applicable

Tunable When disabled

`HADB_MA_START`
The full path to the script that is used to start and stop the HADB Management Agent Server. This script must be able to start and stop the MA Server without any input from the user. Any configuration parameters must be specified in the file indicated by the extension property `HADB_MA_CFG`.

Data type String

Default `/etc/init.d/ma-initd`

Range Not applicable

Tunable When disabled

`HADB_MA_USER`
The user name of the user that starts the HADB Management Agent Server.

Data type String

Default `root`

Range Not applicable

Tunable When disabled

`HADB_ROOT`
The complete path to the HADB installation directory. This directory contains the directory structure `bin/`, which contains the files `ma` and `hadbm`.

Data type String

Default `/opt/SUNWhadb/4`

Range Not applicable

Tunable When disabled

`HADBM_PASSWORDFILE`
The complete path to the file that contains the HADBM administrative password. This property must be set if the HADB management domain is created with an administrative password. If this property is not set, the default null value indicates that no authentication is available for the administrative user.

Data type String

Default Null

Range Not applicable

Tunable At any time

◆ ◆ ◆ APPENDIX B

Properties for Oracle Solaris Cluster HA for Sun Java System Application Server EE (HADB Versions Before 4.4)

This section describes the extension properties for the resource type `SUNW.hadb`. This resource type represents the Sun Java System Application Server EE (HADB) application (supporting HADB versions before 4.4) in a Oracle Solaris Cluster configuration.

For details about system-defined properties, see the `r_properties(5)` man page and the `rg_properties(5)` man page.

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

`Auto_recovery`
Indication of whether to attempt recovery if database is unable to start.

Data type Boolean

Default `FALSE`

Range Not applicable

Tunable At any time

`Auto_recovery_command`
Command to execute after recovering the database. This extension property is optional, regardless of the value of the `Auto_recovery` property.

Data type String

Default Null

Range Not applicable

Tunable At any time

`Confdir_list`
A path name that points to the configuration directory. The data service requires this extension property, and the property must have one entry.

Data type String array

Default No default defined

Range Not applicable

Tunable At creation

`DB_name`
The name of an HADB database. The data service requires this extension property.

Data type String

Default No default defined

Range Not applicable

Tunable When disabled

`DB_password_file`
The file from which HADB reads the password.

Data type String

Default Null

Range Not applicable

Tunable At any time

`HADB_ROOT`
The installation directory location. The data service requires this extension property.

Data type String

Default No default defined

Range Not applicable

Tunable At creation

Index

P

R

S

T

V

Z