Sun Cluster Data Service for Sun Java System HADB Guide for Solaris OS

SPARC Platform Edition

Adobe PostScript™

040129@7940

Contents

Preface

Sun Cluster Data Service for Sun Java System HADB Guide for Solaris OS explains how to install and configure Sun™ Cluster HA for Sun Java System HADB.

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 document assume knowledge of the Solaris™ operating environment and expertise with the volume manager software that is used with Sun Cluster.

UNIX Commands

This document contains information about commands that are specific to installing and configuring Sun 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 Solaris software environment
- Solaris operating environment man pages
- Other software documentation that you received with your system

Typographic Conventions

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

TABLE P–1 Typographic Conventions

Typeface or Symbol	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	Command-line placeholder: replace with a real name or value	To delete a file, type **rm** *filename*.
AaBbCc123	Book titles, new terms, or terms to be emphasized	Read Chapter 6 in *User's Guide*. These are called *class* options. You must be *root* to do this.

Shell Prompts in Command Examples

The following table shows the default system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

TABLE P–2 Shell Prompts

Shell	Prompt
C shell prompt	`machine_name%`
C shell superuser prompt	`machine_name#`
Bourne shell and Korn shell prompt	`$`
Bourne shell and Korn shell superuser prompt	`#`

Related Documentation

Information about related Sun Cluster topics is available in the documentation that is listed in the following table.

Topic	Title	Part Number
Data service administration	*Sun Cluster 3.1 Data Service Planning and Administration Guide*	817-3305
	Sun Cluster 3.1 10/03 Data Services Collection at `http://docs.sun.com/db/coll/573.11`	
Concepts	*Sun Cluster 3.1 10/03 Concepts Guide*	817-0519
Software installation	*Sun Cluster 3.1 10/03 Software Installation Guide*	817-0518
System administration	*Sun Cluster 3.1 10/03 System Administration Guide*	817-0516
Hardware administration	*Sun Cluster 3.1 Hardware Administration Manual*	817-0168
	Sun Cluster 3.x Hardware Administration Collection at `http://docs.sun.com/db/coll/1024.1`	
Data service development	*Sun Cluster 3.1 10/03 Data Services Developer's Guide*	817-0520
Error messages	*Sun Cluster 3.1 10/03 Error Messages Guide*	817-0521
Command and function reference	*Sun Cluster 3.1 10/03 Reference Manual*	817-0522
Release information	*Sun Cluster 3.1 Data Services 10/03 Release Notes*	817-3324
	Sun Cluster 3.1 10/03 Release Notes	817-0638
	Sun Cluster 3.x Release Notes Supplement	816-3381

Accessing Sun Documentation Online

The docs.sun.com^{SM} Web site enables you to access Sun technical documentation online. You can browse the docs.sun.com archive or search for a specific book title or subject. The URL is `http://docs.sun.com`.

Ordering Sun Documentation

Sun Microsystems offers select product documentation in print. For a list of documents and how to order them, see "Buy printed documentation" at http://docs.sun.com.

Help

If you have problems installing or using Sun 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 and serial numbers of your systems
- The release number of the operating environment (for example, Solaris 8)
- The release number of Sun Cluster (for example, Sun Cluster 3.0)

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
scinstall -pv	Displays Sun Cluster release and package version information

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

Installing and Configuring Sun Cluster HA for Sun Java System HADB

This chapter describes the procedures to install and configure the Sun Cluster HA for Sun Java System HADB.

This chapter contains the following procedures.

- "How to Create a Sun Java System HADB Database " on page 13
- "How to Install the Sun Cluster HA for Sun Java System HADB Package by Using the Sun Java Enterprise System Common InstallerProgram " on page 16
- "How to Register and Configure the Sun Cluster HA for Sun Java System HADB" on page 18
- "How to Verify the Sun Cluster HA for Sun Java System HADB Installation and Configuration" on page 23
- "How to Maintain the HADB Database" on page 24

Sun Cluster HA for Sun Java System HADB Overview

This section describes how the Sun Cluster HA for Sun Java System HADB enables the Sun Java System HADB for high availability.

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

Configure the Sun Cluster HA for Sun Java System HADB as a data service mastered by multiple nodes to enable the Sun Java System HADB for high availability. See "Planning for Sun Cluster Data Services" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* for general information about data services.

The Sun Java System HADB software is packaged with the Sun Java System Application Server Enterprise Edition installation. However, you can run Sun Java System 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 HADB in a cluster. First, enable Sun Java System HADB in your cluster to provide session and Enterprise Java Bean (EJB) persistence. Refer to Sun Java System Application Server documentation for information about the Sun Java System HADB. Implementation of the Sun Cluster HA for Sun Java System HADB does not assume the existence of applications on which your architecture depends. However, those applications on which your architecture depends such as databases that are used for invoicing and web servers should be configured to be highly available, might run on a different cluster.

Task Map: Installing and Configuring Sun Cluster HA for Sun Java System HADB

TABLE 1–1 Task Map: Installing and Configuring the Sun Cluster HA for Sun Java System HADB

Task	For Instructions
Plan the Sun Java System HADB installation and configuration	"Planning the Sun Java System HADB Installation and Configuration" on page 11
Install and configure the Sun Java System HADB	"Installing and Configuring the Sun Java System HADB" on page 12
Create a Sun Java System HADB Database	"How to Create a Sun Java System HADB Database " on page 13
Plan the Sun Cluster HA for Sun Java System HADB installation and configuration	"Planning the Sun Cluster HA for Sun Java System HADB Installation and Configuration" on page 14
Install the Sun Cluster HA for Sun Java System HADB package	"Installing the Sun Cluster HA for Sun Java System HADB Package" on page 16
Register and Configure the Sun Cluster HA for Sun Java System HADB as a data service mastered on multiple nodes at one time	"How to Register and Configure the Sun Cluster HA for Sun Java System HADB" on page 18
Configure resource extension properties	"Configuring the Sun Cluster HA for Sun Java System HADB Extension Properties" on page 21

TABLE 1–1 Task Map: Installing and Configuring the Sun Cluster HA for Sun Java System HADB   *(Continued)*

Task	For Instructions
Verify the Sun Cluster HA for Sun Java System HADB installation and configuration	"Verifying the Sun Cluster HA for Sun Java System HADB Installation and Configuration" on page 23
Maintain the HADB database	"Maintaining the HADB Database" on page 23
View fault monitor information	"Operation of the Sun Cluster HA for Sun Java System HADB Fault Monitor" on page 24

Note – If you run multiple data services in your Sun Cluster configuration, you can set up the data services in any order, with the following exception. If the Sun Cluster HA for Sun Java System HADB depends on the Sun Cluster HA for DNS, you must set up DNS first. For details, see *Sun Cluster Data Service for Domain Name Service (DNS) Guide for Solaris OS* . 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 Sun Java System HADB Installation and Configuration

Before you start to install and configure the Sun Java System HADB, use this section in conjunction with the worksheets in "Configuration Worksheets" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

Store static files and data on the local file system of each cluster node. The Sun Java System HADB software is installed when you install the Sun Java System Application Server Enterprise Edition software. See *Sun Cluster Data Service for Sun Java System Application Server Guide for Solaris OS* 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 *Sun Java System Application Server 7 Administrator's Guide*.

Installing and Configuring the Sun Java System HADB

The Sun Java System HADB is a Java 2 Enterprise Edition (J2EE™) 1.3 compliant relational database. The Sun Cluster HA for Sun Java System HADB is designed to meet the needs of enterprise customers and run under the control of Sun 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 *Sun Java System Application Server 7 Administrator's Guide*, the `hadbm` man pages, and the `asadmin` command session-persistence man pages. The Sun Java System HADB is bundled with the Sun Java System Application Server version 7 Enterprise Edition and above. For information about installing and configuring HADB with Sun Java System Application Server, see *Sun Java System Application Server 7 Administrator's Guide*. For information about configuring Sun Cluster HA for Sun Java System Application Server, see *Sun Cluster Data Service for Sun Java System Application Server Guide for Solaris OS*.

Creating a Sun Java System HADB Database

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

- You must specify an even number of Sun Cluster hosts by using the `--hosts` option during database creation.
- You must must specify the Sun Cluster hosts by using the Sun Cluster private interconnect hostnames. An example of a private interconnect hostname is `clusternode1-priv`.
- Sun Java System HADB mirror nodes must be located on different Sun 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 Sun Cluster node is being stopped, the entire database will be shut down. The resource group should only have one node stopped at any time, or all nodes should be stopped with the scswitch -F -g command.
- The HADB resource group must not be used with the Resource Group Offload feature.
- The HADB resource group must not be used with HA Storage Plus.

▼ How to Create a Sun Java System 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 *Sun Java System Application Server 7 Administrator's Guide*.

Note – You must specify the hosts by using the Sun Cluster private interconnect hostnames. To find these hostnames, run the scconf -p | less command from the primary Sun 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 session store and JDBC connection pool. For details, see** *Sun Java System Application Server 7 Administrator's Guide***.**

Planning the Sun Cluster HA for Sun Java System HADB Installation and Configuration

This section contains the information that you need to plan your Sun Cluster HA for Sun Java System HADB installation and configuration.

Configuration Overview

Use the multiple masters configuration in this section to plan the installation and configuration of the Sun Cluster HA for Sun Java System HADB. The Sun Cluster HA for Sun Java System HADB 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 HADB and the Sun Java System Application Server is four active data nodes plus two spare nodes. See "Creating a Sun Java System HADB Database" on page 12 for an example of how to create a Sun Java System HADB with four active 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 Sun Cluster interconnect. See *Sun Java System Application Server 7 Administrator's Guide* in the *Sun Java System Application Server 7 Update 1 Collection (Solaris Edition)* 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 HADB and the Sun Java System Application Server.
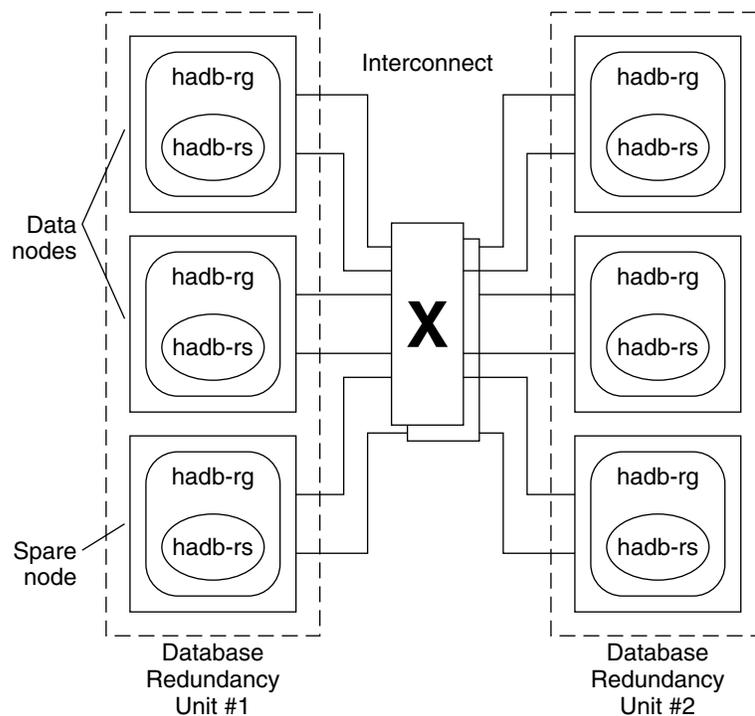
FIGURE 1–1 Recommended Minimum Configuration

Configuration Planning Questions

Use the questions in this section to plan the installation and configuration of the Sun Cluster HA for Sun Java System HADB. See "Considerations" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* for information that might apply to these questions. If you choose to use your HADB data service service with another highly available application, resource dependencies might exist. See "Standard Properties" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* 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 Sun Cluster private interconnect hostnames that you will use?

Installing the Sun Cluster HA for Sun Java System HADB Package

If you did not install the Sun Cluster HA for Sun Java System HADB packages during your initial Sun Cluster installation, perform this procedure to install the packages. Perform this procedure on each cluster node where you are installing the Sun Cluster HA for Sun Java System HADB packages.

If you are installing more than one data service simultaneously, perform the procedure in "Installing the Software" in *Sun Cluster Software Installation Guide for Solaris OS*.

▼ How to Install the Sun Cluster HA for Sun Java System HADB Package by Using the Sun Java Enterprise System Common InstallerProgram

You can run the Sun Java Enterprise System Common 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.

To complete this procedure, you need the Sun Java Enterprise System Common Installer CD-ROM.

1. **On the cluster node where you are installing the Sun Cluster HA for Sun Java System HADB packages, become superuser.**

2. **(Optional) If you intend to run the Sun Java Enterprise System Common Installer program with a GUI, ensure that your `DISPLAY` environment variable is set.**

3. **Load the Sun Java Enterprise System Common Installer CD-ROM into the CD-ROM drive.**

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

4. **Change to the Sun Java Enterprise System Common Installer directory of the CD-ROM.**

   The Sun Java Enterprise System Common Installer resides in this directory.

   ```
   # cd /cdrom/Solaris_sparc
   ```

5. **Start the Sun Java Enterprise System Common Installer program.**

   ```
   # ./installer
   ```

6. **When you are prompted, accept the license agreement and appropriate language support.**

   English language support is available by default.

7. **Select Sun Cluster Agents for Sun Java System under the Availability Services & Sun Cluster 3.1 Subcomponents and proceed.**

   This selection includes all the available Sun Cluster data services for Sun Java System applications, including the Sun Cluster HA for Sun Java System HADB.

8. **When you are prompted, select the type of installation.**

   - If you want minimum configuration performed with the installation, select Custom. You will be prompted for the required configuration information.

   - If you want only to install the data service packages, with no configuration, select Minimal.

9. **(Optional) If you do not want to register the product and receive product updates, uncheck the Product Registration box.**

10. **Follow the instructions on the screen to install the Sun Cluster HA for Sun Java System HADB packages on the node.**

    The Sun Java Enterprise System Common Installer program displays the status of the installation. When the installation is complete, the program displays an installation summary and the installation logs.

11. **Exit the Sun Java Enterprise System Common Installer program.**

    Before exiting the installer program, please make sure that the Sun Cluster HA for Sun Java System HADB has been installed successfully. Check for the presence of the package by executing the following command:

    ```
    # pkginfo -l SUNWschtt
    ```

12. **Unload the Sun Java Enterprise System Common Installer CD-ROM from the CD-ROM drive.**

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

    b. **Eject the CD-ROM.**

       ```
       # eject cdrom
       ```

Registering and Configuring the Sun Cluster HA for Sun Java System HADB

This procedure describes how to use the `scrgadm` command to register and configure the Sun Cluster HA for Sun Java System HADB as a resource mastered on multiple nodes at one time.

Note – See "Tools for Data Service Resource Administration" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* 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 the Sun Cluster HA for Sun Java System HADB which is `SUNW.hadb`
- The names of the cluster nodes that can master the data service

▼ How to Register and Configure the Sun Cluster HA for Sun Java System HADB

Perform the following steps to complete your configuration. For an explanation of the extension properties, see Table 1–2.

1. **Become superuser on a cluster member.**

2. **Register the resource type for the Sun Java System HADB.**

   ```
   # scrgadm -a -t SUNW.hadb
   ```

3. **Create the resource group for the Sun Java System HADB.**

   ```
   # scrgadm -a -g resource_group \
   -y Maximum_primaries=nodes_in_rg \
   -y Desired_primaries=nodes_in_rg \
   -h nodelist
   ```

   `-a`
     Specifies that a resource group is to be created.

   `-g` *resource_group*
     Specifies the name of the resource group that is to be created.

-y 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.

-y 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.

-h *nodelist*
> Specifies a comma-separated subset of cluster nodes to run Sun Java System
> HADB. If this option is omitted, all cluster nodes will run Sun Java System
> HADB. Use the scconf -p | less command to find the nodelist names.

4. Create a Sun Java System HADB resource, either with or without auto-recovery.

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

  ```
  # scrgadm -a -j resource -g resource_group  -t SUNW.hadb \
  -x Confdir_list=config_directory_list \
  -x HADB_ROOT=install_directory \
  -x DB_name=database_name
  ```

  -a
  > Specifies that a data service resource is to be created.

  -j *resource*
  > Specifies that the resource you are creating is named *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.

  -x Confdir_list=*config_directory_list*
  > Specifies the path to the configuration directory.

  -x HADB_ROOT=*install_directory*
  > Specifies the installation directory.

  -x DB_name=*database_name*
  > Specifies the name of the HADB database.

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

  ```
  # scrgadm -a -j resource -g resource_group  -t SUNW.hadb \
  -x Confdir_list=config_directory_list \
  -x HADB_ROOT=install_directory \
  -x DB_name=database_name \
  -x Auto_recovery=true \
  -x Auto_recovery_command=command \
  -x DB_password_file=password_file
  ```

```
-x Auto_recovery=true
```
Specifies that you are using the auto-recovery feature.

```
-x 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.

```
-x DB_password_file=password_file
```
Specifies the file from which HADB will read the database password. See *Sun Java System Application Server 7 Administrator's Guide* for the format and content of the password file.

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.**

    # **scswitch -Z -g** *resource_group*

    -Z   Enables the resource and its monitor

    -g *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.**

    # **scstat -g**
    # **ps -ef**

7. **To verify that you have correctly installed and configured Sun 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.

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

EXAMPLE 1–1 Creating a SUNW.hadb Resource with Auto-recovery

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

In the previous 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`.

Configuring the Sun Cluster HA for Sun Java System HADB Extension Properties

This section describes how to configure the Sun Cluster HA for Sun Java System HADB extension properties. Typically, you use the command line `scrgadm -x` *parameter=value* to configure extension properties when you create the Sun Java System HADB resource. You can also use the procedures described in "Administering Data Service Resources" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* to configure the Sun Cluster HA for Sun Java System HADB extension properties.

See "Standard Properties" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* for details about all Sun Cluster properties.

Table 1–2 describes the extension properties that you can configure for the Sun Java System HADB. The only required extension property for creating a Sun Java System HADB resource is the `Confdir_list` property. The `Confdir_list` property specifies a directory in which the Sun Java System HADB configuration files reside. You can update some extension properties dynamically. You can update other extension properties only when you create the resource. The Tunable entries indicate when you can update each property.

TABLE 1–2 Sun Cluster HA for Sun Java System HADB Extension Properties

Name/Data Type	Description
`Confdir_list` (string array)	A path name that points to the configuration directory. The Sun Cluster HA for Sun Java System HADB requires this extension property, and the property must have one entry. **Default:** `None` **Range:** None **Tunable:** At creation
`Auto_recovery` (boolean)	Attempts recovery if database is unable to start. **Default:** `FALSE` **Tunable:** Any time
`HADB_ROOT` (string)	The installation directory location. The Sun Cluster HA for Sun Java System HADB requires this extension property. **Default:** None **Tunable:** At creation
`Auto_recovery_command` (string)	Command to execute after recovering the database. This extension property is optional, regardless of the value of the `Auto_recovery` property. **Default:** `" "` **Tunable:** Any time
`DB_password_file` (string)	The file from which the password is read by HADB. **Default:** `""` **Tunable:** Any time
`DB_name` (string)	The name of an HADB database. The Sun Cluster HA for Sun Java System HADB requires this extension property. **Default:** None **Tunable:** When disabled

Verifying the Sun Cluster HA for Sun Java System HADB Installation and Configuration

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

▼ How to Verify the Sun Cluster HA for Sun Java System HADB Installation and Configuration

Use this procedure to verify that you installed and configured the Sun Cluster HA for Sun Java System HADB correctly.

1. **Ensure that the HADB is started under the control of Sun Cluster software.**

   # **scswitch -Z -g** *resource_group*

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

   # **scstat -g**
   # **ps -ef**

3. **Verify that you have correctly installed and configured Sun Cluster HA for 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 Sun Cluster HA for Sun Java System HADB data service.

▼ How to Maintain the HADB Database

When you want to run HADB maintenance commands that will 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 once the commands and the rolling restart have completed.

1. **Disable the fault monitor.**

   ```
   scswitch -n -M -j 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.**

   ```
   scswitch -e -M -j resource
   ```

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

Operation of the Sun Cluster HA for Sun Java System HADB Fault Monitor

This section explains the functioning of the Sun Cluster HA for Sun Java System HADB fault monitor.

The start method of the HADB resource starts HADB nodes that are configured to run on the local Sun Cluster node if they are not running. The method then attempts to start the HADB database; in case of failure, the database will be started later in the probe.

The Sun Cluster HA for Sun Java System HADB fault monitor probe periodically checks the status of the HADB database and the HADB nodes. The probe will restart failed HADB nodes and also start the HADB database if the HADB resource was not ready to start the database during the start method. For each iteration of this procedure, the probe executes the following steps:

1. First, the probe sleeps for a period of `Thorough_Probe_Interval` seconds.

2. The probe retrieves the current status of the HADB database and the HADB nodes by executing the `hadbm status` and `hadbm status --nodes` commands.

3. If the database is not running, the probe checks that the HADB `stopstate` file corresponding to that database exists on the local Sun Cluster node. The `hadbm start` command references the `stopstate` file for role assignment of nodes when

it starts the database.

4. 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`.

5. If the database is running, the probe starts the HADB nodes configured to run on the local Sun Cluster node.

6. If the database and the local HADB nodes are running, the probe sets the resource status to `Online` if it was `Online Degraded`.

7. If all the Sun 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.

8. If the `Auto_recovery` extension property has been set to `TRUE`, the HADB resource will attempt to recover the database.

9. If recovery of the database is attempted, the probe executes the following steps:

   - Issues the `hadbm clear --fast` command on one of the Sun Cluster nodes in the resource group's nodelist. 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 Sun 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 take other actions. For example, it may send mail to the Application Server administrator.

   - If both steps succeed, the probe sets the state of the resource to Online.

10. The iteration continues from the first step.

Note – The `Thorough_Probe_Interval` and `Stop_timeout` parameters are tunable with the `scrgadm` command. For details, see "Standard Properties" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

Index

V
verifying, Sun Cluster HA for Sun Java System
    HADB installation, 23