Sun Cluster Data Service for SAP DB Guide for Solaris OS

SPARC Platform Edition

Contents

Preface

Sun Cluster Data Service for SAP DB Guide for Solaris OS explains how to install and configure Sun™ Cluster HA for SAP DB.

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 System 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 Operating System
- Solaris Operating System man pages
- Other software documentation that you received with your system

Typographic Conventions

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

TABLE P–1 Typographic Conventions

Typeface or Symbol	Meaning	Example
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your `.login` file. Use `ls -a` to list all files. `machine_name% you have mail.`
AaBbCc123	What you type, contrasted with on-screen 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 words, or terms, or words 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. All Sun Cluster documentation is available at `http://docs.sun.com`.

Topic	Documentation
Data service administration	*Sun Cluster Data Services Planning and Administration Guide for Solaris OS*
	Individual data service guides
Concepts	*Sun Cluster Concepts Guide for Solaris OS*
Overview	*Sun Cluster Overview for Solaris OS*
Software installation	*Sun Cluster Software Installation Guide for Solaris OS*
System administration	*Sun Cluster System Administration Guide for Solaris OS*
Hardware administration	*Sun Cluster 3.x Hardware Administration Manual for Solaris OS*
	Individual hardware administration guides
Data service development	*Sun Cluster Data Services Developer's Guide for Solaris OS*
Error messages	*Sun Cluster Error Messages Guide for Solaris OS*
Command and function reference	*Sun Cluster Reference Manual for Solaris OS*

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

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 Solaris Operating System (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
SPARC: `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 SAP DB

This chapter explains how to install and configure Sun Cluster HA for SAP DB.

This chapter contains the following procedures.

- "How to Install and Configure SAP DB" on page 17
- "How to Enable SAP DB to Run in a Cluster" on page 18
- "How to Verify SAP DB Installation and Configuration on Each Node" on page 20
- "How to Install the Sun Cluster HA for SAP DB Packages by Using the Web Start Program" on page 22
- "How to Install the Sun Cluster HA for SAP DB Packages by Using the `scinstall` Utility" on page 24
- "How to Register and Configure an `HAStoragePlus` Resource" on page 25
- "How to Register and Configure an SAP xserver Resource" on page 28
- "How to Register and Configure an SAP DB Resource" on page 30
- "How to Verify the Operation of the SAP DB Fault Monitor" on page 38
- "How to Verify the Operation of the SAP xserver Fault Monitor" on page 40

Sun Cluster HA for SAP DB Overview

To eliminate single points of failure in an SAP DB system, Sun Cluster HA for SAP DB provides the following features:

- Fault monitoring and automatic failover for the SAP DB application. You must configure Sun Cluster HA for SAP DB as a failover data service.
- Fault monitoring and automatic restart for SAP xserver. You must configure SAP xserver as a scalable data service.

For conceptual information about failover data services and scalable data services, see the *Sun Cluster Concepts Guide for Solaris OS*.

Each component of SAP DB has data service that protects the component when the component is configured in Sun Cluster. See the following table.

Note – The files that are associated with the `SUNW.sap_xserver` resource type are supplied with the Sun Cluster HA for SAP liveCache data service. The Sun Cluster HA for SAP liveCache data service is installed when you install Sun Cluster HA for SAP DB data service.

TABLE 1–1 Protection of SAP DB Components by Sun Cluster Data Services

SAP DB Component	Data Service
SAP DB	Sun Cluster HA for SAP DB
	The resource type is `SUNW.sapdb`.
SAP xserver	Sun Cluster HA for SAP liveCache
	The resource type is `SUNW.sap_xserver`.
NFS file system	Sun Cluster HA for NFS
	The resource type is `SUNW.nfs`.
	For more information about this data service, see *Sun Cluster Data Service for Network File System (NFS) Guide for Solaris OS*.

Overview of the Installation and Configuration Process for Sun Cluster HA for SAP DB

The following table summarizes the tasks for installing and configuring Sun Cluster HA for SAP DB 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 1–2 Tasks for Installing and Configuring Sun Cluster HA for SAP DB

Task	Cross-Reference
Plan the Sun Cluster HA for SAP DB installation and configuration	Your SAP documentation. *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* "Planning the Sun Cluster HA for SAP DB Installation and Configuration" on page 12
Install and configure SAP DB	"Installing and Configuring SAP DB" on page 16
Verify the SAP DB installation and configuration	"Verifying the SAP DB Installation and Configuration" on page 20
Install the Sun Cluster HA for SAP DB packages	"Installing the Sun Cluster HA for SAP DB Packages" on page 21
Configure the `HAStoragePlus` resource to work with Sun Cluster HA for SAP DB	"Relationship Between Resource Groups and Disk Device Groups" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* "Synchronizing the Startups Between Resource Groups and Disk Device Groups" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* "Configuring the `HAStoragePlus` Resource Type to Work With Sun Cluster HA for SAP DB" on page 24
Register and configure the Sun Cluster HA for SAP DB data service	"Registering and Configuring Sun Cluster HA for SAP DB" on page 26
(Optional) Tune the Sun Cluster HA for SAP DB fault monitors	"Tuning the Sun Cluster HA for SAP DB Fault Monitors" on page 32
(Optional) Use an alternate project identifier (ID)	"Cluster Administration and Application Development" in *Sun Cluster Concepts Guide for Solaris OS*
Verify the Sun Cluster HA for SAP DB installation and configuration	"Verifying the Sun Cluster HA for SAP DB Installation and Configuration" on page 38
(Optional) Upgrade the `SUNW.sap_xserver` resource type	"Upgrading the `SUNW.sap_xserver` Resource Type" on page 41

Planning the Sun Cluster HA for SAP DB Installation and Configuration

This section contains the information that you need to plan your Sun Cluster HA for SAP DB installation and configuration.

Note – Before you begin, consult your SAP DB documentation for configuration restrictions and requirements that are not imposed by Sun Cluster software. For information about restrictions that the Sun Cluster software imposes, see the Sun Cluster documentation.

Configuration Requirements

The configuration requirements in this section apply only to Sun Cluster HA for SAP DB.

Caution – If your data service configuration does not conform to these requirements, the data service configuration might not be supported.

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

SAP DB Software Version Requirements

Use SAP DB version 7.4 or compatible versions.

Sun Cluster HA for SAP DB Configuration Requirements

Configure Sun Cluster HA for SAP DB as a failover data service. You cannot configure Sun Cluster HA for SAP DB as a scalable data service. For more information, see the following sections:

- "How to Enable SAP DB to Run in a Cluster" on page 18
- "How to Register and Configure an SAP DB Resource" on page 30

SAP xserver Configuration Requirements

To enable client applications to access Sun Cluster HA for SAP DB, you must use SAP xserver. Configure SAP xserver as a scalable data service. Do *not* configure SAP xserver as a failover data service.

Configure SAP xserver so that SAP xserver starts on all nodes to which the SAP DB resource can fail over. To implement this configuration, ensure that the node list of the SAP xserver resource group contains all nodes that are in the node list of the SAP DB resource group. For more information, see "How to Register and Configure an SAP xserver Resource" on page 28.

Supported Configurations of This Data Service

The Sun Cluster HA for SAP DB data service supports configurations that conform to the requirements in "Configuration Requirements" on page 12.

If you plan to use SAP DB with other highly available SAP applications, you must also configure the Sun Cluster data services for those applications. For more information, see the following table.

SAP Application	Sun Cluster Data Service	Associated Document
SAP R/3	Sun Cluster HA for SAP	*Sun Cluster Data Service for SAP Guide for Solaris OS*
SAP liveCache	Sun Cluster HA for SAP liveCache	*Sun Cluster Data Service for SAP liveCache Guide for Solaris OS*

The examples that follow show these supported configurations of Sun Cluster HA for SAP DB:

- Two-node cluster configuration
- Four-node cluster configuration with SAP R/3
- Four-node cluster configuration with SAP R/3 and SAP liveCache

---

Note – Sun Cluster HA for SAP DB might support additional configurations. However, you must contact your Sun service provider for information about additional configurations.

EXAMPLE 1–1 Two-Node Configuration

This example shows a two-node configuration in which a client application accesses the SAP DB resource through the SAP xserver resource. The characteristics of this configuration are as follows:

EXAMPLE 1–1 Two-Node Configuration　　*(Continued)*

- The SAP DB resource is configured as a failover data service.
- The SAP xserver resource is configured as a scalable data service.

```
┌──────────────┐   ┌──────────────┐
│  ╭────────╮  │   │ ┌──────────┐ │
│  │  SAP   │  │   │ │  Client  │ │
│  │  DB    │  │   │ │application│ │
│  ╰────────╯  │   │ └──────────┘ │
│  ╭────────╮  │   │ ╭──────────╮ │
│  │  SAP   │──┼───┼─│   SAP    │ │
│  │ xserver│  │   │ │ xserver  │ │
│  ╰────────╯  │   │ ╰──────────╯ │
└──────────────┘   └──────────────┘
```

EXAMPLE 1–2 Four-Node Configuration With SAP R/3

This example shows a four-node configuration in which SAP DB is used with SAP R/3. This configuration uses multiple Advanced Planner & Optimizer (APO) application servers. The characteristics of this configuration are as follows:

- The SAP APO Central Instance (CI) resource is configured as a failover data service.
- The SAP DB resource is configured as a failover data service.
- The SAP xserver resource is configured as a scalable data service.
- APO application server (APP) resources are configured as scalable data services.

```
┌──────────┐  ┌──────────┐  ┌──────────┐  ┌──────────┐
│ ╭──────╮ │  │ ╭──────╮ │  │ ╭──────╮ │  │ ╭──────╮ │
│ │  CI  │ │  │ │ SAP  │ │  │ │ APP  │ │  │ │ APP  │ │
│ ╰──────╯ │  │ │  DB  │ │  │ ╰──────╯ │  │ ╰──────╯ │
│ ╭──────╮ │──│ ╰──────╯ │──│ ╭──────╮ │──│ ╭──────╮ │
│ │ APP  │ │  │ ╭──────╮ │  │ │ SAP  │ │  │ │ SAP  │ │
│ ╰──────╯ │  │ │ SAP  │ │  │ │xserver│ │  │ │xserver│ │
│          │  │ │xserver│ │  │ ╰──────╯ │  │ ╰──────╯ │
└──────────┘  │ ╰──────╯ │  └──────────┘  └──────────┘
              └──────────┘
```

EXAMPLE 1–3 Four-Node Configuration With SAP R/3 and SAP liveCache

This example shows a four-node configuration in which SAP DB is used with SAP R/3 and SAP liveCache. This configuration uses multiple APO application servers. The characteristics of this configuration are as follows:

- The SAP APO CI resource is configured as a failover data service.
- The SAP DB resource is configured as a failover data service.
- The SAP xserver resource is configured as a scalable data service.
- APP resources are configured as scalable data services.

■ The SAP liveCache resource is configured as a failover data service.

Configuration Considerations

The configuration considerations in the subsections that follow affect the installation and configuration of Sun Cluster HA for SAP DB.

Device Group for the SAP DB Application

Ensure that you create a device group for the SAP DB application as follows:

■ Install SAP DB on its own global device group. For more information, see "Installing and Configuring SAP DB" on page 16. This separate global device group for SAP DB ensures that the SAP DB resource can depend only on the `HAStoragePlus` resource for SAP DB.

■ Create an `HAStoragePlus` resource for the global device group on which SAP DB is installed. For more information, see "Configuring the `HAStoragePlus` Resource Type to Work With Sun Cluster HA for SAP DB" on page 24.

■ Ensure that the resource for Sun Cluster HA for SAP DB depends on the `HAStoragePlus` resource for the global device group on which SAP DB is installed. For more information, see "Registering and Configuring Sun Cluster HA for SAP DB" on page 26.

Administration of SAP xserver by a User Other Than Root

You might be required to administer SAP xserver as a user other than root. In this situation, you must create and define that user as follows:

■ You must create that user on all cluster nodes that master SAP xserver.

■ You must define that user when you register and configure Sun Cluster HA for SAP DB. For more information, see "Registering and Configuring Sun Cluster HA for SAP DB" on page 26.

Configuration Planning Questions

Answer the questions in this section to plan the installation and configuration of Sun Cluster HA for SAP DB. Write the answers to these questions in the space that is provided on the data service worksheets in "Configuration Worksheets" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

■ Which resource group will you use for the SAP DB application resource and for the logical hostname for the SAP DB resource?

Use the answer to this question when you perform the procedure "How to Enable SAP DB to Run in a Cluster" on page 18.

■ What is the logical hostname for the SAP DB resource? Clients access the data service through this logical hostname.

Use the answer to this question when you perform the following procedures:

■ "How to Install and Configure SAP DB" on page 17
■ "How to Enable SAP DB to Run in a Cluster" on page 18

■ Where will the system configuration files reside?

See *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* for the advantages and disadvantages of using the local file system instead of the cluster file system.

Installing and Configuring SAP DB

To enable Sun Cluster HA for SAP DB to make SAP DB highly available, additional installation and configuration operations are required. These operations supplement the standard installation and standard configuration of SAP DB.

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

For information about the standard installation and standard configuration of SAP DB, see the following documentation:

■ If you are using SAP DB with SAP R/3, see the SAP R/3 documentation for information about how to install and configure SAP R/3 with SAP DB.

■ If you are using SAP DB independently of SAP R/3, see the SAP DB documentation.

▼ How to Install and Configure SAP DB

1. **On one node of the cluster, install the SAP DB software.**

   Ensure that you install SAP DB on its own global device group.

2. **Perform a standard configuration of SAP DB.**

3. **Create the `.XUSER.62` file in the home directory of the operating system (OS) user who administers the SAP DB instance.**

   a. **Create a plain text file that contains information about the database user who administers the SAP DB database instance.**

      For information about the content of this file, see the SAP DB documentation. For the name of the server on which the database is running, specify the logical hostname for the SAP DB resource that you specified in "Configuration Planning Questions" on page 16.

      For an example of the content of this file, see Example 1–4.

   b. **Generate the `.XUSER.62` file from the plain text file that you created in Step a.**

      Use the SAP DB command `xuser` for this purpose.

      # **xuser -b** *user-info-file*

      `-b` *user-info-file*   Specifies the plain text file from which you are generating the `.XUSER.62` file

4. **Copy the `/usr/spool/sql` directory and its contents from the node on which you installed SAP DB to all nodes where resources for SAP DB and SAP xserver will run.**

   To ensure that the same owner owns the directory and its contents on all nodes, use the `tar`(1) command and the `rsh`(1) command.

   # **tar cfB - /usr/spool/sql | rsh** *destination* **tar xfB -**

   destination   Specifies the node to which you are copying the `/usr/spool/sql` directory and its contents

The following example shows a plain text file that contains information about a database user who administers an SAP DB instance.

EXAMPLE 1–4 Information About a Database User Who Administers an SAP DB Instance

```
DEFAULT
dbm
dbm
TST
srvr-1
    blank line
    blank line
-1
```

EXAMPLE 1–4 Information About a Database User Who Administers an SAP DB Instance *(Continued)*

```
-1
```
   blank line

This example specifies the following information about a database user who administers an SAP DB instance:

- The user key that is used to address this combination of XUSER data is named DEFAULT.
- The user name and password of the database user are dbm.
- The name of the SAP DB instance is TST.
- The logical hostname for the SAP DB resource is srvr-1.
- No structured query language (SQL) mode is specified.
- The default time-out value of the SAP DB instance is used.
- The default isolation level of the SAP DB instance is used.

For more information, see the SAP DB documentation.

▼ How to Enable SAP DB to Run in a Cluster

1. **Create a failover resource group to contain the SAP DB application resources and the logical hostname for the SAP DB resource.**

   Use the resource group that you identified when you answered the questions in "Configuration Planning Questions" on page 16.

   ```
   # scrgadm -a -g sapdb-rg [ -h nodelist]
   ```

   -a   Specifies that a new configuration is to be added.

   -g *sapdb-rg*   Specifies that the resource group that you are creating is named *sapdb-rg*.

   -h *nodelist*   Specifies a comma-separated list of nodes where the resource group can be brought online. Specifying a node list is optional. If you omit the node list, the resource group can be brought online on all cluster nodes.

2. **Ensure that all network resources that you intend to you use are added to your name service database.**

3. **Add a logical hostname resource to the failover resource group that you created in Step 1.**

   ```
   # scrgadm -a -L -g sapdb-rg -l sapdb-logical-hostname [ -n netiflist]
   ```

-a

    Specifies that a new configuration is to be added.

-L

    Species that a logical hostname resource is to be added.

-g *sapdb-rg*

    Specifies that the logical hostname resource is to be added to the failover
    resource group that you created in Step 1.

-l *sapdb-logical-hostname*

    Specifies the logical hostname of the server on which the database is running.
    This hostname must be the logical hostname for the SAP DB resource that you
    specified in "Configuration Planning Questions" on page 16.

-n *netiflist*

    Specifies a comma-separated list of network interfaces. All nodes in the node list
    of the resource group must be represented in the list of network interfaces.
    Specifying a list of network interfaces is optional. If you omit the list, an attempt
    is made for each node to discover a network interface on the subnet that the
    hostname list identifies.

4. **Enable the resource group that you created in Step 1.**

    # **scswitch -Z -g** *sapdb-rg*

    -Z   Moves a resource group to the MANAGED state and brings the
                    resource group online

    -g *sapdb-rg*   Specifies that the resource group that you created in Step 1 is to be
                    moved to the MANAGED state and brought online

EXAMPLE 1–5 Enabling SAP DB to Run in a Cluster

This example shows the sequence of commands that are required to enable SAP DB to
run in a cluster. The commands are run on only one cluster node.

1. The following command creates a failover resource group to contain the SAP DB
   application resources and the logical hostname for the SAP DB resource. The
   resource group is named sapdbrg. The sapdbrg resource group can be brought
   online on all cluster nodes.

   # **scrgadm -a -g sapdbrg**

2. The following command adds a logical hostname resource to the sapdbrg resource
   group. The logical hostname of the server on which the database is running is
   srvr-1. When the sapdbrg resource group is brought online, an attempt is made
   for each node to discover a network interface on the subnet that the hostname list
   identifies.

   # **scrgadm -a -L -g sapdbrg -l srvr-1**

3. The following command moves the sapdbrg resource group to the MANAGED state
   and brings the resource group online.

EXAMPLE 1–5 Enabling SAP DB to Run in a Cluster   *(Continued)*

```
# scswitch -Z -g sapdbrg
```

Verifying the SAP DB Installation and Configuration

Before you install the Sun Cluster HA for SAP DB packages, verify that the SAP DB software is correctly installed and configured to run in a cluster. This verification does *not* verify that the SAP DB application is highly available because the Sun Cluster HA for SAP DB data service is not yet installed.

▼ How to Verify SAP DB Installation and Configuration on Each Node

Perform this procedure on each node that can master the SAP DB resource group.

1. **Log in as superuser to a node that can master the SAP DB resource group.**

2. **Switch the SAP DB resource group to the node that you logged in to in Step 1.**

   ```
   # scswitch -z -g sapdb-rg -h node
   ```

   -z   Specifies that the node that masters a resource group is to be switched

   -g *sapdb-rg*   Specifies that the SAP DB resource group is to be switched to another node

   -h *node*   Specifies the node to which the SAP DB resource group is to be switched

3. **Confirm that the SAP DB database can be started and be stopped.**

   a. **Become the OS user who administers the SAP DB database.**

      ```
      # su - os-sapdb-adm-user
      ```

      os-sapdb-adm-user   Specifies the UNIX user identity of the OS user who administers the SAP DB database. This user's home directory contains the `.XUSER.62` file that was created during the installation and configuration of SAP DB. For more information, see "Installing and Configuring SAP

DB" on page 16. You specify this user when you set the
`DB_User` extension property as explained in "How to
Register and Configure an SAP DB Resource" on page 30.

b. Start the SAP xserver.

```
# x_server start
```

c. Manually start the SAP DB database instance on the node that you logged in to in Step 1.

```
# dbmcli -U sapdb-adm-key db_online
```

`-U` *sapdb-adm-key*   Specifies that the `dbmcli` command is run with the user key of the database user who administers the SAP DB instance. This user key is created during the installation and configuration of SAP DB. For more information, see "Installing and Configuring SAP DB" on page 16. You specify this user key when you set the `User_Key` extension property as explained in "How to Register and Configure an SAP DB Resource" on page 30.

d. Confirm that the SAP DB database instance is started.

e. Manually stop the SAP DB database instance.

```
# dbmcli -U sapdb-adm-key db_offline
```

`-U` *sapdb-adm-key*   Specifies that the `dbmcli` command is run with the user key that you used for starting the database in Step c

f. Confirm that the SAP DB database instance is stopped.

Installing the Sun Cluster HA for SAP DB Packages

If you did not install the Sun Cluster HA for SAP DB 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 SAP DB packages. To complete this procedure, you need the Sun Java Enterprise System Accessory CD Volume 3.

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*.

Install the Sun Cluster HA for SAP DB packages by using one of the following installation tools:

- The Web Start program
- The scinstall utility

Note – The Web Start program is *not* available in releases earlier than Sun Cluster 3.1 Data Services 10/03.

Requirements for Installing With an Existing Version of Sun Cluster HA for SAP liveCache

The files that are associated with the SUNW.sap_xserver resource type are supplied with the Sun Cluster HA for SAP liveCache data service. The installation tools install the Sun Cluster HA for SAP liveCache data service when you install the Sun Cluster HA for SAP DB data service.

If you are using any version of Sun Cluster HA for SAP liveCache in the following list, you must upgrade to version 3.1 4/04 of Sun Cluster HA for SAP liveCache *before* you install the Sun Cluster HA for SAP DB data service:

- Version 3.0 5/02 asynchronous release
- Version 3.1 5/03
- Version 3.1 10/03

For instructions for upgrading Sun Cluster software, see "Upgrading Sun Cluster Software" in *Sun Cluster Software Installation Guide for Solaris OS*.

▼ How to Install the Sun Cluster HA for SAP DB Packages by Using the Web Start Program

You can run the Web Start 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. For more information about the Web Start program, see the installer(1M) man page.

1. **On the cluster node where you are installing the Sun Cluster HA for SAP DB packages, become superuser.**

2. **(Optional) If you intend to run the Web Start program with a GUI, ensure that your DISPLAY environment variable is set.**

3. **Load the Sun Java Enterprise System Accessory CD Volume 3 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/cdrom0` directory.

4. **Change to the Sun Cluster HA for SAP DB component directory of the CD-ROM.**

   The Web Start program for the Sun Cluster HA for SAP DB data service resides in this directory.

   ```
   # cd /cdrom/cdrom0/\
   components/SunCluster_HA_SAPDB_1.0
   ```

5. **Start the Web Start program.**

   ```
   # ./installer
   ```

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

   - To install only the C locale, select Typical.
   - To install other locales, select Custom.

7. **Follow instructions on the screen to install the Sun Cluster HA for SAP DB packages on the node.**

   After the installation is finished, the Web Start program provides an installation summary. This summary enables you to view logs that the Web Start program created during the installation. These logs are located in the `/var/sadm/install/logs` directory.

8. **Exit the Web Start program.**

9. **Unload the Sun Java Enterprise System Accessory CD Volume 3 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
      ```

Where to Go From Here

Go to "Configuring the `HAStoragePlus` Resource Type to Work With Sun Cluster HA for SAP DB" on page 24.

▼ How to Install the Sun Cluster HA for SAP DB Packages by Using the `scinstall` Utility

1. **Load the Sun Java Enterprise System Accessory CD Volume 3 into the CD-ROM drive.**

2. **Run the `scinstall` utility with no options.**

   This step starts the `scinstall` utility in interactive mode.

3. **Select the menu option, Add Support for New Data Service to This Cluster Node.**

   The `scinstall` utility prompts you for additional information.

4. **Provide the path to the Sun Java Enterprise System Accessory CD Volume 3.**

   The utility refers to the CD as the "data services cd."

5. **Specify the data service to install.**

   The `scinstall` utility lists the data service that you selected and asks you to confirm your choice.

6. **Exit the `scinstall` utility.**

7. **Unload the CD from the drive.**

Where to Go From Here

Configuring the `HAStoragePlus` Resource Type to Work With Sun Cluster HA for SAP DB

For maximum availability of the SAP DB database, resources that Sun Cluster HA for SAP DB requires must be available before the SAP DB database instance is started. An example of such a resource is the file system where programs and libraries for the SAP DB runtime environment reside. To ensure that these resources are available, configure the `HAStoragePlus` resource type to work with Sun Cluster HA for SAP DB.

For information about the relationship between resource groups and disk device groups, see "Relationship Between Resource Groups and Disk Device Groups" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

Configuring the `HAStoragePlus` resource type to work with Sun Cluster HA for SAP DB involves the following operations:

- Synchronizing the startups between resource groups and disk device groups as explained in "Synchronizing the Startups Between Resource Groups and Disk Device Groups" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*

- Registering and configuring an `HAStoragePlus` resource

▼ How to Register and Configure an `HAStoragePlus` Resource

1. **Register the `SUNW.HAStoragePlus` resource type.**

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

2. **Create an `HAStoragePlus` resource for the global device group on which SAP DB is installed.**

   Create this resource in the SAP DB resource group. This resource must perform an affinity switchover for all global devices that are defined for this resource.

   ```
   # scrgadm -a -j hsp-resource -g sapdb-rg \
   -t SUNW.HAStoragePlus -x filesystemmountpoints=mountpoint-list \
   -x globaldevicepaths=sapdb-device-group -x affinityon=TRUE
   ```

   `-a`
   : Specifies that a new configuration is to be added.

   `-j` *hsp-resource*
   : Specifies that the resource that you are creating is named *hsp-resource*.

   `-g` *sapdb-rg*
   : Specifies that the resource is to be added to the SAP DB resource group.

   `-t SUNW.HAStoragePlus`
   : Specifies that the resource is an instance of the `SUNW.HAStoragePlus` resource type.

   `-x filesystemmountpoints=`*mountpoint-list*
   : Specifies a list of valid mount points for the file system. For more information, see the `SUNW.HAStoragePlus`(5) man page.

   `-x globaldevicepaths=`*sapdb-device-group*
   : Specifies the name of the global device group on which the SAP DB software is installed.

```
-x affinityon=TRUE
```
   Specifies that this resource performs an affinity switchover for all global devices
   that are defined for this resource.

3. **Enable the `HAStoragePlus` resource that you created in Step 2.**

   ```
   # scswitch -ej hsp-resource
   ```

   `-ej` *hsp-resource*   Specifies that the `HAStoragePlus` resource that you created
            in Step 2 is to be enabled

EXAMPLE 1–6 Creating an `HAStoragePlus` Resource

```
# scrgadm -a -j hsprs  -g sapdbrg \
-t SUNW.HAStoragePlus -x filesystemmountpoints=/global/sapdbdata \
-x globaldevicepaths=sapdbdg  -x affinityon=TRUE
```

This example shows the creation of a `SUNW.HAStoragePlus` resource that has the
following characteristics:

- The resource is named `hsprs`.
- The resource is a member of a resource group that is named `sapdbrg`. The creation
  of this resource group is shown in Example 1–5.
- The resource is an instance of the `SUNW.HAStoragePlus` resource type. The
  registration of this resource type is not shown in this example.
- The mount point for the file system is `/global/sapdbdata`.
- The SAP DB software is installed on a global device group that is named `sapdbdg`.
- The `hsprs` resource performs an affinity switchover for all global devices that are
  defined for this resource.

Where to Go From Here

Go to "Registering and Configuring Sun Cluster HA for SAP DB" on page 26.

Registering and Configuring Sun Cluster HA for SAP DB

To enable Sun Cluster HA for SAP DB to make SAP DB highly available, configure Sun
Cluster data services as follows:

- Configure SAP xserver as a scalable data service.
- Configure Sun Cluster HA for SAP DB as a failover data service.

Before you perform this procedure, ensure that the Sun Cluster HA for SAP DB data service packages are installed.

⚠ **Caution –** One SAP xserver serves multiple SAP DB instances and, if SAP liveCache is used, multiple SAP liveCache instances in the cluster. Therefore, do *not* configure more than one SAP xserver resource on the same cluster. If more than one SAP xserver resource runs on the same cluster, conflicts between the SAP xserver resources occur. These conflicts cause all SAP xserver resources to become unavailable. If you attempt to start the SAP xserver a second time, the attempt fails. The error message `Address already in use` is also displayed.

Setting Sun Cluster HA for SAP DB Extension Properties

The sections that follow contain instructions for registering and configuring resources. These instructions explain how to set *only* extension properties that Sun Cluster HA for SAP DB requires you to set. For information about all Sun Cluster HA for SAP DB extension properties, see Appendix A. You can update some extension properties dynamically. You can update other properties, however, only when you create or disable a resource. The Tunable entry indicates when you can update a property.

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

-x *property=value*

-x *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 "Administering Data Service Resources" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* to configure resources after the resources are created.

Administering SAP xserver as a User Other Than Root

You might be required to administer SAP xserver as a user other than root. In this situation, you must create and define that user as follows:

- You must create that user on all cluster nodes that master SAP xserver.

- You must define that user when you register and configure an SAP xserver resource. To define the user who administers SAP xserver, set the `Xserver_User` extension property when you create an SAP xserver resource. For more information about the `Xserver_User` extension property, see "`SUNW.sap_xserver` Extension Properties" on page 46.

▼ How to Register and Configure an SAP xserver Resource

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

2. **Register the `SUNW.sap_xserver` resource type.**

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

3. **Create a scalable resource group for the SAP xserver resource.**

   Configure SAP xserver so that SAP xserver starts on all nodes to which the SAP DB resource can fail over. To implement this configuration, ensure that the node list of the SAP xserver resource group contains all nodes that are in the node list of the SAP DB resource group. This resource group is created when the procedure "How to Enable SAP DB to Run in a Cluster" on page 18 is performed.

   ```
   # scrgadm -a -g xserver-rg \
   -y Maximum_primaries=nodes-in-sapdb-rg \
   -y Desired_primaries=nodes-in-sapdb-rg \
   -h nodelist
   ```

   `-a`
   Specifies that a new configuration is to be added.

   `-g` *xserver-rg*
   Specifies that the resource group that you are creating is named *xserver-rg*.

   `-y Maximum_primaries=`*nodes-in-sapdb-rg*
   Specifies the maximum number of nodes on which the SAP xserver resource can start. This number is equal to the number of nodes that are in the node list of the SAP DB resource group. You must specify the same number as the value of the `Desired_primaries` property.

   `-y Desired_primaries=`*nodes-in-sapdb-rg*
   Specifies the desired number of nodes on which the SAP xserver resource can start. This number is equal to the number of nodes that are in the node list of the SAP DB resource group. You must specify the same number as the value of the `Maximum_primaries` property.

   `-h` *nodelist*
   Specifies a comma-separated list of nodes where the resource group can be brought online. Ensure that this node list contains all nodes that are in the node list of the SAP DB resource group.

4. **Create an SAP xserver resource in the resource group that you created in Step 3.**

   # **scrgadm -a -j** *xserver-resource* **-g** *xserver-rg* **-t SUNW.sap_xserver**

   -a
     Specifies that a new configuration is to be added

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

   -g *xserver-rg*
     Specifies that the resource is to be added to the resource group that you created
     in Step 3

   -t SUNW.sap_xserver
     Specifies that the resource is an instance of the SUNW.sap_xserver resource
     type

5. **Enable the resource group that you created in Step 3.**

   # **scswitch -Z -g** *xserver-rg*

   -Z   Moves a resource group to the MANAGED state and brings the
                     resource group online

   -g *xserver-rg*   Specifies that the resource group that you created in Step 3 is to be
                     moved to the MANAGED state and brought online

6. **Modify the SAP DB resource group to depend on the resource group that you
   created in Step 3.**

   # **scrgadm -c -g** *sapdb-rg* **-y rg_dependencies=***xserver-rg*

   -c
     Specifies that an existing configuration is to be modified

   -g *sapdb-rg*
     Specifies that the SAP DB resource group is to be modified

   -y rg_dependencies=*xserver-rg*
     Specifies that the SAP DB resource group is to depend on the resource group
     that you created in Step 3

EXAMPLE 1–7 Configuring a SUNW.sap_xserver Resource

This example shows the sequence of commands that are required to configure an SAP
xserver resource. The commands are run on only one cluster node.

1. The following command creates a scalable resource group to contain an SAP
   xserver resource for a four-node cluster. The resource group is named xsrvrrg.
   The xsrvrrg resource group can be brought online on all cluster nodes.

   ```
   # scrgadm -a -g xsrvrrg \
   -y Maximum_primaries=4 \
   -y Desired_primaries=4
   ```

EXAMPLE 1–7 Configuring a `SUNW.sap_xserver` Resource   *(Continued)*

2. The following command creates an SAP xserver resource that is named `xsrvrrs` in the `xsrvrrg` resource group. The SAP xserver resource is an instance of the `SUNW.sap_xserver` resource type. The registration of this resource type is not shown in this example.

   ```
   # scrgadm -a -j xsrvrrs  -g xsrvrrg -t SUNW.sap_xserver
   ```

3. The following command moves the `xsrvrrg` resource group to the `MANAGED` state and brings the resource group online.

   ```
   # scswitch -Z -g
   ```

4. The following command modifies the `sapdbrg` resource group to depend on the `xsrvrrg` resource group. The creation of the `sapdbrg` resource group is shown in Example 1–5.

   ```
   # scrgadm -c -g sapdbrg -y rg_dependencies=xsrvrrg
   ```

▼ How to Register and Configure an SAP DB Resource

1. **Register the `SUNW.sapdb` resource type.**

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

2. **Create an SAP DB resource in the SAP DB resource group.**

   Ensure that the SAP DB resource depends on the `HAStoragePlus` resource for the global device group on which SAP DB is installed.

   When you create this resource, specify the following information about the SAP DB database instance. This information is created when SAP DB is installed and configured as explained in "Installing and Configuring SAP DB" on page 16.

   - The name of the SAP DB database
   - The UNIX user identity of the OS user who administers the SAP DB database instance
   - The user key of the database user who administers the SAP DB database instance

   ```
   # scrgadm -a -j sapdb-rs -g sapdb-rg -t SUNW.sapdb \
   -x DB_Name=db-name -x DB_User=os-sapdb-adm-user \
   -x User_Key=sapdb-adm-key -y resource_dependencies=hsp-resource
   ```

   -a
     Specifies that a new configuration is to be added.

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

-g *sapdb-rg*
: Specifies that the resource is to be added to the SAP DB resource group.

-t SUNW.sapdb
: Specifies that the resource is an instance of the SUNW.sapdb resource type.

-x DB_Name=*db-name*
: Specifies the name of the SAP DB database instance in uppercase.

-x DB_User=*os-sapdb-adm-user*
: Specifies the UNIX user identity of the OS user who administers the SAP DB database. This user's home directory contains the .XUSER.62 file that was created during the installation and configuration of SAP DB. For more information, see "Installing and Configuring SAP DB" on page 16.

-x User_Key=*sapdb-adm-key*
: Specifies the user key of the database user who administers the SAP DB database instance. This user key is created during the installation and configuration of SAP DB. For more information, see "Installing and Configuring SAP DB" on page 16.

-y resource_dependencies=*hsp-resource*
: Specifies that the SAP DB resource depends on the HAStoragePlus resource for the global device group on which SAP DB is installed.

3. **Enable the SAP DB resource group.**

   ```
   # scswitch -Z -g sapdb-rg
   ```

   -Z
   : Moves a resource group to the MANAGED state and brings the resource group online

   -g *sapdb-rg*
   : Specifies that the SAP DB resource group is to be moved to the MANAGED state and brought online

4. **(Optional) Consider configuring your cluster to offload noncritical resource groups.**

   You might plan to run lower-priority services on a node to which the SAP DB resource can fail over. In this situation, consider using an RGOffload resource to shut down those services when the SAP DB resource fails over to the node. For more information, see "Freeing Node Resources by Offloading Non-critical Resource Groups" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

EXAMPLE 1–8 Creating a SUNW.sapdb Resource

```
# scrgadm -a -j sapdbrs -g sapdbrg -t SUNW.sapdb \
-x DB_Name=TST -x DB_User=dbadmin \
-x User_Key=DEFAULT -y resource_dependencies=hsprs
```

This example shows the creation of a SUNW.sapdb resource that has the following characteristics:

EXAMPLE 1–8 Creating a `SUNW.sapdb` Resource   *(Continued)*

- The resource is named `sapdbrs`.
- The resource is a member of a resource group that is named `sapdbrg`. The creation of this resource group is shown in Example 1–5.
- The resource is an instance of the `SUNW.sapdb` resource type. The registration of this resource type is not shown in this example.
- The SAP DB database instance that is associated with this resource is named `TST`.
- The UNIX user identity of the OS user who administers the SAP DB database is `dbadmin`.
- The user key of the database user who administers the SAP DB database is `DEFAULT`.
- The SAP DB resource depends on an `HAStoragePlus` resource that is named `hsprs`. The creation of the `hsprs` resource is shown in Example 1–6.

Tuning the Sun Cluster HA for SAP DB Fault Monitors

Fault monitoring for the Sun Cluster HA for SAP DB data service is provided by the following fault monitors:

- The SAP DB fault monitor
- The SAP xserver fault monitor

Each fault monitor is contained in a resource whose resource type is shown in the following table.

TABLE 1–3 Resource Types for Sun Cluster HA for SAP DB Fault Monitors

Fault Monitor	Resource Type
SAP DB	`SUNW.sapdb`
SAP xserver	`SUNW.sap_xserver`

System properties and extension properties of these resources control the behavior of the fault monitors. The default values of these properties determine the preset behavior of the fault monitors. The preset behavior should be suitable for most Sun Cluster installations. Therefore, you should tune the Sun Cluster HA for SAP DB fault monitors *only* if you need to modify this preset behavior.

Tuning the Sun Cluster HA for SAP DB fault monitors involves the following tasks:

- Setting the interval between fault monitor probes
- Setting the timeout for fault monitor probes
- Specifying the response to a fault

Perform these tasks when you register and configure Sun Cluster HA for SAP DB. For more information, see "Registering and Configuring Sun Cluster HA for SAP DB" on page 26.

Operation of the Sun Cluster HA for SAP DB Fault Monitors

To determine whether SAP xserver and the SAP DB database instance are operating correctly, the Sun Cluster HA for SAP DB fault monitors probe these resources periodically.

During a probe, the SAP DB fault monitor performs the following operations:

1. The SAP DB fault monitor determines whether the SAP DB database instance is online.

2. If the SAP DB database instance is online, the SAP DB fault monitor determines whether the parent kernel process of the SAP DB database instance is running. You can control how the fault monitor responds if the parent kernel process is not running. For more information, see "Forcing the SAP DB Database Instance to Be Restarted if the Parent Kernel Process Is Terminated" on page 37.

3. The SAP DB fault monitor determines whether SAP xserver is available. This fault monitoring supplements the fault monitoring that the SAP xserver fault monitor provides.

During a probe, the SAP xserver fault monitor determines whether SAP xserver is available.

The SAP DB fault monitor responds to the following faults in SAP DB:

- A status of the SAP DB database instance that is not `ONLINE`, for example, `OFFLINE` or `ADMIN`
- Unexpected termination of the parent kernel process of the SAP DB database instance

The SAP DB fault monitor also responds to the unavailability of SAP xserver. This fault monitoring supplements the fault monitoring that the SAP xserver fault monitor provides.

The SAP xserver fault monitor responds to following faults:

- Unavailability of SAP xserver
- Persistent system errors

Setting the Interval Between Fault Monitor Probes

To determine whether SAP xserver and the SAP DB database instance are operating correctly, the Sun Cluster HA for SAP DB fault monitors probe these resources periodically. The interval between fault monitor probes affects the availability of each resource and the performance of your system as follows:

- The interval between fault monitor probes affects the length of time that is required to detect a fault and respond to the fault. Therefore, if you decrease the interval between fault monitor probes, the time that is required to detect a fault and respond to the fault is also decreased. This decrease enhances the availability of the resource.
- Each fault monitor probe consumes system resources such as processor cycles and memory. Therefore, if you decrease the interval between fault monitor probes, the performance of the system is degraded.

The optimum interval between fault monitor probes also depends on the time that is required to respond to a fault in the resource. This time depends on how the complexity of the resource affects the time that is required for operations such as restarting the resource.

For example, SAP xserver is a much simpler resource and can be restarted much quicker than SAP DB. Therefore, the optimum interval between fault monitor probes of SAP xserver is shorter than the optimum interval between probes of SAP DB.

To set the interval between fault monitor probes, set the `Thorough_probe_interval` system property to the interval in seconds that you require. Set this property for each resource that contains a Sun Cluster HA for SAP DB fault monitor that you need to tune. The resource types of these resources are shown in Table 1–3.

Setting the Timeout for Fault Monitor Probes

The timeout for fault monitor probes specifies the length of time that a fault monitor waits for a response from a resource to a probe. If the fault monitor does not receive a response within this timeout, the fault monitor treats the resource as faulty. The time that a resource requires to respond to a fault monitor probe depends on the operations that the fault monitor performs to probe the resource.

During a probe, the SAP DB fault monitor performs the following operations:

1. The SAP DB fault monitor determines whether the SAP DB database instance is online.
2. If the SAP DB database instance is online, the SAP DB fault monitor determines whether the parent kernel process of the SAP DB database instance is running.

3. The SAP DB fault monitor determines whether SAP xserver is available. This fault monitoring supplements the fault monitoring that the SAP xserver fault monitor provides.

During a probe, the SAP xserver fault monitor determines whether SAP xserver is available.

The time that is required for a resource to respond also depends on factors that are unrelated to the fault monitor or the application, for example:

- System configuration
- Cluster configuration
- System load
- Amount of network traffic

To set the timeout for fault monitor probes, set the `Probe_timeout` extension property to the timeout in seconds that you require. Set this property for each resource that contains a Sun Cluster HA for SAP DB fault monitor that you need to tune. The resource types of these resources are shown in Table 1–3.

Specifying the Response to a Fault

The SAP DB fault monitor responds to the following faults in SAP DB:

- A status of the SAP DB database instance that is not `ONLINE`, for example, `OFFLINE` or `ADMIN`
- Unexpected termination of the parent kernel process of the SAP DB database instance

The SAP DB fault monitor also responds to the unavailability of SAP xserver. This fault monitoring supplements the fault monitoring that the SAP xserver fault monitor provides.

The SAP xserver fault monitor responds to following faults:

- Unavailability of SAP xserver
- Persistent system errors

The Sun Cluster HA for SAP DB fault monitors enable you to specify the response to these faults. Specifying the response to a fault involves the following tasks:

- Defining the criteria for persistent faults

- Disabling the failover of the SAP DB resource

- Forcing the SAP DB database instance to be restarted if the parent kernel process is terminated

Defining the Criteria for Persistent Faults

To minimize the disruption that transient faults in a resource cause, a fault monitor restarts the resource in response to such faults. For persistent faults, more disruptive action than restarting the resource is required:

- For the SAP DB resource, the fault monitor fails over the resource to another node. The SAP DB resource is a failover resource.
- For the SAP xserver resource, the fault monitor takes the resource offline. The SAP xserver is a scalable resource.

The fault monitors treat a fault as persistent if the number of attempts to restart a resource exceeds a specified threshold within a specified retry interval. Defining the criteria for persistent faults enables you to set the threshold and the retry interval to accommodate the performance characteristics of your cluster and your availability requirements.

Dependencies of the Threshold and the Retry Interval on Other Properties

The maximum length of time that is required for a single restart of a faulty resource is the sum of the values of the following properties:

- `Thorough_probe_interval` system property
- `Probe_timeout` extension property

To ensure that you allow enough time for the threshold to be reached within the retry interval, use the following expression to calculate values for the retry interval and the threshold:

$$\textit{retry-interval} \geq \textit{threshold} \times (\textit{thorough-probe-interval} + \textit{probe-timeout})$$

System Properties for Setting the Threshold and the Retry Interval

To set the threshold and the retry interval, set the following system properties:

- To set the threshold, set the `Retry_count` system property to the maximum allowed number of restarts.
- To set the retry interval, set the `Retry_interval` system property to the interval in seconds that you require.

Set these properties for each resource that contains a Sun Cluster HA for SAP DB fault monitor that you need to tune. The resource types of these resources are shown in Table 1–3.

Other Effects of the Retry Interval

Besides defining a criterion for persistent faults, the retry interval affects the response of a fault monitor to the following faults:

- **Unavailability of SAP xserver that the** *SAP DB* **fault monitor detects.** If the SAP DB fault monitor detects that SAP xserver is unavailable twice within the retry interval, the SAP DB fault monitor restarts SAP xserver.

- **Persistent system errors.** A persistent system error is a system error that occurs four times within the retry interval. If a persistent system error occurs, the fault monitor restarts SAP xserver.

Disabling the Failover of the SAP DB Resource

By default, failover of the SAP DB resource is enabled. The fault monitor fails over the SAP DB resource in response to a persistent fault. A persistent fault occurs if the number of attempts to restart exceeds `Retry_count` within the time that `Retry_interval` specifies.

Disable the failover of the SAP DB resource if you require the fault monitor to restart the SAP DB even in response to a persistent fault. When failover is disabled, the fault monitor reports a persistent fault in a resource by setting the status of the resource to faulted.

To disable the failover of the SAP DB resource, set the `Failover_enabled` extension property of the `SUNW.sapdb` resource to `False`.

Forcing the SAP DB Database Instance to Be Restarted if the Parent Kernel Process Is Terminated

By default, unexpected termination of the parent kernel process does *not* cause the SAP DB fault monitor to restart the SAP DB database instance. The SAP DB database instance can continue to function without the parent kernel process. Restarting the SAP DB database instance in this situation might cause unnecessary unavailability of the SAP DB database instance. Therefore, you should force the SAP DB database instance to be restarted *only* if you require a feature that the parent kernel process provides. An example of such a feature is maintaining the integrity of the log history.

To force the SAP DB database instance to be restarted if the parent kernel process is terminated, set the `Restart_if_Parent_Terminated` extension property of the `SUNW.sapdb` resource to `True`.

Verifying the Sun Cluster HA for SAP DB Installation and Configuration

After you install, register, and configure Sun Cluster HA for SAP DB, verify the Sun Cluster HA for SAP DB installation and configuration. Verifying the Sun Cluster HA for SAP DB installation and configuration determines if the Sun Cluster HA for SAP DB data service makes the SAP DB application highly available.

Verifying the Sun Cluster HA for SAP DB installation involves verifying the operation of the following fault monitors:

- The SAP DB fault monitor
- The SAP xserver fault monitor

▼ How to Verify the Operation of the SAP DB Fault Monitor

Perform this procedure on each node where SAP DB can run.

1. **Log in as superuser to a node that can master the SAP DB resource group.**

2. **Switch the SAP DB resource group to the node that you logged in to in Step 1.**

   # **scswitch -z -g** *sapdb-rg* **-h** *node*

   -z   Specifies that the node that masters a resource group is to be switched

   -g *sapdb-rg*   Specifies that the SAP DB resource group is to be switched to another node

   -h *node*   Specifies the node to which the SAP DB resource group is to be switched

3. **Abnormally terminate SAP DB.**

   a. **Determine the process identities of all kernel processes for the SAP DB database instance that you are running.**

      # **ps -ef | grep kernel | grep** *db-name*

      db-name   Specifies the name of the SAP DB database instance in uppercase

   b. **Kill all kernel processes for the SAP DB database instance that you are running.**

      # **kill -9** *sapdb-kernel-pid*

sapdb-kernel-pid Specifies the process identities of the SAP DB kernel processes that you determined in Step a

4. **Confirm that the Sun Cluster HA for SAP DB fault monitor performs the appropriate operation from the following list:**

   - Restarting the SAP DB resource
   - Failing over the SAP DB resource to another node

   The expected behavior of the fault monitor depends on the failure history of the resource and the value of the `Failover_enabled` extension property. For more information, see "Defining the Criteria for Persistent Faults" on page 36 and "Forcing the SAP DB Database Instance to Be Restarted if the Parent Kernel Process Is Terminated" on page 37.

5. **Terminate SAP DB normally.**

   a. **Become the OS user who administers the SAP DB database.**

      ```
      # su - os-sapdb-adm-user
      ```

      os-sapdb-adm-user Specifies the UNIX user identity of the OS user who administers the SAP DB database. This user's home directory contains the `.XUSER.62` file that was created during the installation and configuration of SAP DB. For more information, see "Installing and Configuring SAP DB" on page 16. You specify this user when you set the `DB_User` extension property as explained in "How to Register and Configure an SAP DB Resource" on page 30.

   b. **Manually stop the SAP DB database instance.**

      ```
      # dbmcli -U sapdb-adm-key db_offline
      ```

      `-U` *sapdb-adm-key* Specifies that the `dbmcli` command is run with the user key of the database user who administers the SAP DB database. This user key is created during the installation and configuration of SAP DB. For more information, see "Installing and Configuring SAP DB" on page 16. You specify this user key when you set the `User_Key` extension property as explained in "How to Register and Configure an SAP DB Resource" on page 30.

6. **Confirm that the Sun Cluster HA for SAP DB fault monitor performs the appropriate operation from the following list:**

   - Restarting the SAP DB resource
   - Failing over the SAP DB resource to another node

   The expected behavior of the fault monitor depends on the failure history of the resource and the value of the `Failover_enabled` extension property. For more information, see "Defining the Criteria for Persistent Faults" on page 36 and

"Forcing the SAP DB Database Instance to Be Restarted if the Parent Kernel Process Is Terminated" on page 37.

▼ How to Verify the Operation of the SAP xserver Fault Monitor

Perform this procedure on each node that can master SAP xserver.

1. **Log in to a node that can master SAP xserver.**

2. **Abnormally terminate SAP xserver.**

   a. **Determine the process identities of all SAP xserver processes.**

      # **ps -ef | grep vserver**

   b. **Kill all SAP xserver processes.**

      # **kill -9** *xserver-pid*

      xserver-pid   Specifies the process identities of the SAP xserver processes that you determined in Step a

3. **Confirm that the SAP xserver fault monitor restarts the SAP xserver resource.**

4. **Terminate SAP xserver normally.**

   a. **Become the OS user who administers SAP xserver.**

      # **su -** *os-sapxsrvr-adm-user*

      os-sapxsrvr-adm-user
      Specifies the UNIX user identity of the OS user who administers SAP xserver. By default, this user is root. You can specify this user by setting the Xserver_User extension property. For more information, see "SUNW.sap_xserver Extension Properties" on page 46.

   b. **Manually stop the SAP xserver.**

      # **x_server stop**

5. **Confirm that the SAP xserver fault monitor restarts the SAP xserver resource.**

Upgrading the `SUNW.sap_xserver` Resource Type

The `SUNW.sap_xserver` resource type is supplied with the Sun Cluster HA for SAP liveCache data service. The Sun Cluster HA for SAP liveCache data service is installed when you install Sun Cluster HA for SAP DB data service. Upgrade the `SUNW.sap_xserver` resource type if all conditions in the following list apply:

- You are using version 3.0 5/02 asynchronous release, version 3.1 5/03, or version 3.1 10/03 of the Sun Cluster HA for SAP liveCache data service.

- You plan to use the Sun Cluster HA for SAP DB data service with your existing version of the Sun Cluster HA for SAP liveCache data service.

- You need to use the new features of the `SUNW.sap_xserver` resource type.

For general instructions that explain how to upgrade a resource type, see "Upgrading a Resource Type" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*. The information that you need to complete the upgrade of the `SUNW.sap_xserver` resource type is provided in the subsections that follow.

Information for Registering the New Resource Type Version

The relationship between a resource type version and the release of Sun Cluster data services is shown in the following table. The release of Sun Cluster data services indicates the release in which the version of the resource type was introduced.

Resource Type Version	Sun Cluster Data Services Release
1.0	3.0 5/02 asynchronous release
2	3.1 4/04

To determine the version of the resource type that is registered, use one command from the following list:

- `scrgadm -p`
- `scrgadm -pv`

The resource type registration (RTR) file for this resource type is `/opt/SUNWsclc/xserver/etc/SUNW.sap_xserver`.

Information for Migrating Existing Instances of the Resource Type

The information that you need to migrate instances of the `SUNW.sap_xserver` resource type is as follows:

- You can perform the migration only when the resource is unmonitored.
- If you need to use the new features of the `SUNW.sap_xserver` resource type, the required value of the `Type_version` property is 2.
- If you need to specify the directory that contains programs and libraries for the SAP xserver runtime environment, set the `Independent_Program_Path` extension property. For more information, see "`SUNW.sap_xserver` Extension Properties" on page 46.

The following example shows a command for editing an instance of the `SUNW.sap_xserver` resource type.

EXAMPLE 1–9 Editing an Instance of the `SUNW.sap_xserver` Resource Type During Upgrade

```
# scrgadm -cj sapxserver-rs -y Type_version=2 \
  -x Independent_Program_Path=/sapdb/indep_prog
```

This command edits a `SUNW.sap_xserver` resource as follows:

- The `SUNW.sap_xserver` resource is named `sapxserver-rs`.
- The `Type_version` property of this resource is set to 2.
- The independent program path is `/sapdb/indep_prog`.

Sun Cluster HA for SAP DB Extension Properties

Extension properties for Sun Cluster HA for SAP DB resource types are described in the following sections.

- "SUNW.sapdb Extension Properties" on page 43
- "SUNW.sap_xserver Extension Properties" on page 46

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

SUNW.sapdb Extension Properties

The SUNW.sapdb resource type represents the SAP DB application in a Sun Cluster configuration. The extension properties of this resource type are as follows:

dbmcli_Start_Option
The option that is passed to the dbmcli command to start the SAP DB database instance.

Note – For SAP DB version 7.4.3, set this property to db_online.

Data type	String
Default	db_online
Range	Not applicable
Tunable	When disabled

DB_Name
   The name of the SAP DB database instance in uppercase. This name is created
   when SAP DB is installed and configured as explained in "Installing and
   Configuring SAP DB" on page 16.

Data type	String
Default	No default defined
Range	Not applicable
Tunable	When disabled

DB_User
   The UNIX user identity of the operating system (OS) user who administers the SAP
   DB database instance. This user's home directory contains the .XUSER.62 file that
   was created during the installation and configuration of SAP DB. For more
   information, see "Installing and Configuring SAP DB" on page 16.

Data type	String
Default	No default defined
Range	Not applicable
Tunable	When disabled

Failover_enabled
   Specifies whether the fault monitor fails over the SAP DB resource if the number of
   attempts to restart exceeds Retry_count within the time that Retry_interval
   specifies. The possible values of this extension property are as follows:

   - True – Specifies that the fault monitor fails over the SAP DB resource
   - False – Specifies that the fault monitor does *not* fail over the SAP DB resource

Data type	Boolean
Default	True
Range	Not applicable
Tunable	Any time

Independent_Program_Path
   The full path to the directory that contains the following programs and libraries for
   the SAP DB application:

   - Programs that are independent of the database software version
   - Libraries for the client runtime environment

   Sun Cluster HA for SAP DB determines the path to the dbmcli command from the
   value of this property. The dbmcli command resides in the bin subdirectory of the
   directory that this property specifies.

Data type	String
Default	/sapdb/programs

Range   Not applicable

Tunable   When disabled

Monitor_retry_count
The maximum number of restarts by the process monitor facility (PMF) that are allowed for the fault monitor.

Data type   Integer

Default   4

Range   No range defined

Tunable   Any time

Monitor_retry_interval
The period of time in minutes during which the PMF counts restarts of the fault monitor.

Data type   Integer

Default   2

Range   No range defined

Tunable   Any time

Pid_Dir_Path
The full path to the directory under which files that store the process identities of SAP DB kernel processes are created. The process identities of SAP DB kernel processes are stored in the following files:

- *pid-dir*/ppid/*db-name*
- *pid-dir*/pid/*db-name*

The replaceable items in these file paths are as follows:

- *pid-dir* is the directory that the Pid_Dir_Path extension property specifies
- *db-name* is the name of the SAP DB database instance that the DB_Name extension property specifies

Data type   String

Default   /var/spool/sql

Range   Not applicable

Tunable   When disabled

Probe_timeout
The time-out value in seconds that the fault monitor uses to probe an SAP DB database instance.

Data type   Integer

Default   90

Range   30–99,999

Tunable   Any time

`Restart_if_Parent_Terminated`
Determines whether the fault monitor restarts the SAP DB database instance if the parent kernel process is terminated. The possible values of this extension property are as follows:

- `True` – Specifies that the fault monitor restarts the SAP DB database instance if the parent kernel process is terminated
- `False` – Specifies that the fault monitor does *not* restart the SAP DB database instance if the parent kernel process is terminated

Data type   Boolean

Default   `False`

Range   Not applicable

Tunable   Any time

`User_Key`
The user key of the database user who administers the SAP DB database instance. This user key is created when SAP DB is installed and configured as explained in "Installing and Configuring SAP DB" on page 16.

Data type   String

Default   No default defined

Range   Not applicable

Tunable   When disabled

`SUNW.sap_xserver` Extension Properties

The `SUNW.sap_xserver` resource type represents SAP xserver in a Sun Cluster configuration. The extension properties of this resource type are as follows:

`Confdir_List`
The full path to the directory that contains the SAP DB software and SAP DB database instance.

Data type   String

Default   `/sapdb`

Range   Not applicable

Tunable   At creation

`Independent_Program_Path`
  The full path to the directory that contains the following programs and libraries for
  SAP xserver:

- Programs that are independent of the database software version
- Libraries for the client runtime environment

Sun Cluster HA for SAP DB determines the path to the `x_server` command from
the value of this property. The `x_server` command resides in the `bin` subdirectory
of the directory that this property specifies.

Data type	String
Default	No default defined
Range	Not applicable
Tunable	When disabled
Introduced in release	3.1 4/04

`Monitor_retry_count`
  The maximum number of restarts by the PMF that are allowed for the fault
  monitor.

Data type	Integer
Default	4
Range	No range defined
Tunable	Any time

`Monitor_retry_interval`
  The period of time in minutes during which the PMF counts restarts of the fault
  monitor.

Data type	Integer
Default	2
Range	No range defined
Tunable	Any time

`Probe_timeout`
  The time-out value in seconds for fault monitor probes.

Data type	Integer
Default	120
Range	No range defined
Tunable	Any time

Soft_Stop_Pct

The percentage of the Stop method timeout that is used to stop SAP xserver by using the SAP utility x_server stop. If this timeout is exceeded, the SIGKILL signal is used to stop all SAP xserver processes.

Data type	Integer
Default	50
Range	1–100
Tunable	When disabled

Xserver_User

The UNIX user identity of the OS user who administers SAP xserver.

Data type	String
Default	root
Range	Not applicable
Tunable	At creation

Index