Sun Cluster 3.1 Data Service for Sybase ASE

Adobe PostScript™

030314@5533

Contents

Preface

The Sun Cluster 3.1 Data Service for Sybase ASE contains procedures to install and configure the Sun Cluster data services.

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 used with Sun Cluster.

UNIX Commands

This document contains information on commands specific to installing and configuring Sun Cluster data services. It might not contain information on basic UNIX® commands and procedures, such as shutting down the system, booting the system, and configuring devices. For that information, see one or more of the following:

- Online documentation for the Solaris software environment
- Solaris operating environment man pages
- Other software documentation that you received with your system

Typographic Conventions

Typeface or Symbol	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your file. Use `ls -a` to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% **su** Password:
AaBbCc123	Book titles, new words or terms, words to be emphasized	Read Chapter 6 in the *User's Guide*. These are called *class* options. You *must* be superuser to do this.
	Command-line variable; replace with a real name or value	To delete a file, type `rm` *filename*.

Shell Prompts

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

Related Documentation

Application	Title	Part Number
Installation	*Sun Cluster 3.1 Software Installation Guide*	816-3388
Data Services Installation and Configuration	*Sun Cluster 3.1 Data Service Planning and Administration Guide*	817–1526
	Sun Cluster 3.1 Data Service for Apache	817–1532
	Sun Cluster 3.1 Data Service for BroadVision One-To-One Enterprise	817–1542
	Sun Cluster 3.1 Data Service for Domain Name Service (DNS)	817–1533
	Sun Cluster 3.1 Data Service for Netbackup	817–1538
	Sun Cluster 3.1 Data Service for Network File System (NFS)	817–1534
	Sun Cluster 3.1 Data Service for Oracle	817–1527
	Sun Cluster 3.1 Data Service for Oracle Parallel Server/Real Application Clusters	817–1535
	Sun Cluster 3.1 Data Service for SAP	817–1536
	Sun Cluster 3.1 Data Service for SAP liveCache	817–1539
	Sun Cluster 3.1 Data Service for Siebel	817–1540
	Sun Cluster 3.1 Data Service for Sun ONE Application Server	817–1530
	Sun Cluster 3.1 Data Service for Sun ONE Directory Server	817–1529
	Sun Cluster 3.1 Data Service for Sun ONE Message Queue	817–1531
	Sun Cluster 3.1 Data Service for Sun ONE Web Server	817–1528
	Sun Cluster 3.1 Data Service for Sybase ASE	817–1537
	Sun Cluster 3.1 Data Service for WebLogic Server	817–1537
	Sun Cluster 3.1 Data Service for DHCP	817–1716
	Sun Cluster 3.1 Data Service for Samba	817–1715
	Sun Cluster 3.1 Data Service for WebSphere MQ	817–1714

Application	Title	Part Number
	Sun Cluster 3.1 Data Service for WebSphere MQ Integrator	817–1713
Hardware	*Sun Cluster 3.x Hardware Administration Manual*	817–0168
	Sun Cluster 3.x Data Service Collection at `http://docs.sun.com/db/coll/1024.1`	
API development	*Sun Cluster 3.1 Data Services Developer's Guide*	816-3385
Administration	*Sun Cluster 3.1 5/03 System Administration Guide*	816-3384
Cluster concepts	*Sun Cluster 3.1 Concepts Guide*	816-3383
Error Messages	*Sun Cluster 3.1 Error Messages Guide*	816-3382
Man Pages	*Sun Cluster 3.1 Reference Manual*	816–5251
Release notes	*Sun Cluster 3.1 Data Service 5/03 Release Notes*	817–1790
	Sun Cluster 3.1 Release Notes	816-5317
	Sun Cluster 3.1 Release Notes Supplement	816-3381

Sun Documentation Online

The `docs.sun.com`^{SM} Web site enables you to access Sun technical documentation on the Web. You can browse the `docs.sun.com` archive or search for a specific book title or subject 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 E-mail 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 Sybase ASE

This chapter provides instructions on how to configure and administer Sun Cluster HA for Sybase ASE on your Sun Cluster nodes.

This chapter contains the following procedures.

- "How to Prepare the Nodes" on page 13
- "How to Install the Sybase ASE Software" on page 15
- "How to Verify the Sybase ASE Installation" on page 17
- "How to Configure Sybase ASE Database Access With Solstice DiskSuite/Solaris Volume Manager" on page 17
- "How to Configure Sybase ASE Database Access With VERITAS Volume Manager" on page 18
- "How to Create the Sybase ASE Database Environment" on page 19
- "How to Install Sun Cluster HA for Sybase ASE Packages" on page 21
- "How to Register and Configure Sun Cluster HA for Sybase ASE" on page 22
- "How to Verify the Sun Cluster HA for Sybase ASE Installation " on page 25

You must configure Sun Cluster HA for Sybase ASE as a failover data service. See the *Sun Cluster 3.1 Concepts Guide* document and "Planning for Sun Cluster Data Services" in *Sun Cluster 3.1 Data Service Planning and Administration Guide* for general information about data services, resource groups, resources, and other related topics.

Installing and Configuring Sun Cluster HA for Sybase ASE

The following table lists sections that describe the installation and configuration tasks.

TABLE 1–1 Task Map: Installing and Configuring Sun Cluster HA for Sybase ASE

Task	For Instructions, Go To
Prepare to install Sun Cluster HA for Sybase ASE	"Preparing to Install Sun Cluster HA for Sybase ASE" on page 12
Install the Sybase ASE 12.0 software	"Installing the Sybase ASE 12.0 Software" on page 13
Create the Sybase database environment	"Creating the Sybase ASE Database Environment" on page 17
Install the Sun Cluster HA for Sybase ASE package	"Installing the Sun Cluster HA for Sybase ASE Package" on page 21
Register Sun Cluster HA for Sybase ASE resource types and configure resource groups and resources	"Registering and Configuring Sun Cluster HA for Sybase ASE" on page 22
Verify the Sun Cluster HA for Sybase ASE installation	"Verifying the Sun Cluster HA for Sybase ASE Installation" on page 25
Understand Sun Cluster HA for Sybase ASE logging and security issues	"Understanding Sun Cluster HA for Sybase ASE Logging and Security Issues" on page 26
Configure Sun Cluster HA for Sybase ASE extension properties	"Configuring Sun Cluster HA for Sybase ASE Extension Properties" on page 27
View fault monitor information	"Sun Cluster HA for Sybase ASE Fault Monitor" on page 31

Preparing to Install Sun Cluster HA for Sybase ASE

To prepare your nodes for the Sun Cluster HA for Sybase Adaptive Server 12.0 installation, select an installation location for the following files.

■ **Sybase ASE application files –** These files include Sybase ASE binaries and libraries. You can install these files on either the local file system or the cluster file system.

See "Configuration Guidelines for Sun Cluster Data Services" in *Sun Cluster 3.1 Data Service Planning and Administration Guide* for the advantages and disadvantages of placing the Sybase ASE binaries on the local file system as opposed to the cluster file system.

■ **Sybase ASE configuration files –** These files include the `interfaces` file, `config` file, and environment file. You can install these files on the local file system (with links), the highly available local file system, or on the cluster file system.

■ **Database data files –** These files include Sybase device files. You must install these files on the highly available local file system or the cluster file system as either raw devices or regular files.

Installing the Sybase ASE 12.0 Software

Use the procedures in this section to complete the following tasks.

■ Prepare the nodes.
■ Install the Sybase ASE software.
■ Verify the Sybase ASE installation.

Note – Before you configure Sun Cluster HA for Sybase ASE, use the procedures that the *Sun Cluster 3.1 Software Installation Guide* describes to configure the Sun Cluster software on each node.

▼ How to Prepare the Nodes

This procedure describes how to prepare the cluster nodes for Sybase ASE software installation.

Caution – Perform all of the steps in this procedure on all of the nodes. If you do not perform all of the steps on all of the nodes, the Sybase ASE installation will be incomplete, and Sun Cluster HA for Sybase ASE will fail during startup.

Note – Consult the Sybase ASE documentation before you perform this procedure.

1. **Become superuser on all of the nodes.**

2. **Configure the `/etc/nsswitch.conf` file as follows so that Sun Cluster HA for Sybase ASE starts and stops correctly if a switchover or failover occurs.**

   On each node that can master the logical host that runs Sun Cluster HA for Sybase ASE, include one of the following entries for group in the /etc/nsswitch.conf file.

   ```
   group:
   group: files [NOTFOUND=return] nis
   ```

```
group: file [NOTFOUND=return] nisplus
```

Sun Cluster HA for Sybase ASE uses the su *user* command to start and stop the database node.

The network information name service might become unavailable when a cluster node's public network fails. Adding one of the preceding entries for `group` ensures that the su(1M) command does not refer to the NIS/NIS+ name services if the network information name service is unavailable.

3. **Configure the cluster file system for Sun Cluster HA for Sybase ASE.**

   If raw devices contain the databases, configure the global devices for raw-device access. See the *Sun Cluster 3.1 Software Installation Guide* for information on how to configure global devices.

   If you use the Solstice DiskSuite/Solaris Volume Manager, configure the Sybase ASE software to use UNIX file system (UFS) logging on mirrored meta devices or raw-mirrored meta devices. See the Solstice DiskSuite/Solaris Volume Manager documentation for information on how to configure raw-mirrored metadevices.

4. **Prepare the SYBASE_HOME directory on a local or multihost disk.**

   Note – If you install the Sybase ASE binaries on a local disk, use a separate disk if possible. Installing the Sybase ASE binaries on a separate disk prevents the binaries from overwrites during operating environment reinstallation.

   ---

5. **On each node, create an entry for the database administrator (DBA) group in the /etc/group file, and add potential users to the group.**

   Verify that the `root` and *sybase* users are members of the *dba* group, and add entries as necessary for other DBA users. Ensure that group IDs are the same on all of the nodes that run Sun Cluster HA for Sybase ASE, as the following example illustrates.

   dba`:*:520:root`,*sybase*

   You can create group entries in a network name service. If you do so, also add your entries to the local /etc/group file to eliminate dependency on the network name service.

6. **On each node, create an entry for the Sybase system administrator.**

   The following command updates the /etc/passwd and /etc/shadow files with an entry for the Sybase system administrator.

   # **useradd -u 120 -g** *dba* **-d /**Sybase-home *sybase*

   Ensure that the *sybase* user entry is the same on all of the nodes that run Sun Cluster HA for Sybase ASE.

▼ How to Install the Sybase ASE Software

Perform the following steps to install the Sybase ASE software.

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

2. **Note the Sybase ASE installation requirements.**
   You can install Sybase ASE binaries on one of the following locations.

   - Local disks of the cluster nodes
   - Highly available local file system
   - Cluster file system

   > **Note –** Before you install the Sybase ASE software on the cluster file system, start the Sun Cluster software and become the owner of the disk device group.

   See "Preparing to Install Sun Cluster HA for Sybase ASE" on page 12 for more information about installation locations.

3. **Create a failover resource group to hold the network and application resources.**

   ```
   # scrgadm -a -g resource-group [-h nodelist]
   ```

   -g *resource-group*   Specifies the name of the resource group. This name can be your choice but must be unique for resource groups within the cluster.

   -h *nodelist*   Specifies an optional, comma-separated list of physical node names or IDs that identify potential masters. The order here determines the order in which the Resource Group Manager (RGM) considers primary nodes during failover.

   > **Note –** Use the -h option to specify the order of the node list. If all of the nodes in the cluster are potential masters, you do not need to use the -h option.

4. **Verify that you have added all of the network resources that Sun Cluster HA for Sybase ASE uses to either the /etc/inet/hosts file or to your name service (NIS, NIS+) database.**

5. **Add a network resource (logical hostname or shared address) to the failover resource group.**

   ```
   # scrgadm -a -L -g resource-group -l logical-hostname [-n netiflist]
   ```

-l *logical-hostname*	Specifies a network resource. The network resource is the logical hostname or shared address (IP address) that clients use to access Sun Cluster HA for Sybase ASE.
-n *netiflist*	Specifies an optional, comma-separated list that identifies the IP Networking Multipathing groups that are on each node. Each element in *netiflist* must be in the form of `netif@node`. `netif` can be given as an IP Networking Multipathing group name, such as `sc_ipmp0`. The node can be identified by the node name or node ID, such as `sc_ipmp0@1` or `sc_ipmp@phys-schost-1`.

Note – Sun Cluster does not currently support using the adapter name for `netif`.

---

6. **Run the `scswitch(1M)` command to complete the following tasks.**

   - Enable the resource and fault monitoring.
   - Move the resource group into a managed state.
   - Bring the resource group online.

   ```
   # scswitch -Z -g resource-group
   ```

7. **On the node mastering the resource group that you just created, login as** *sybase*.

   The installation of the Sybase binaries must be performed on the node where the corresponding logical host is running.

8. **Install the Sybase ASE software.**

   Regardless of where you install the Sybase ASE software, modify each node's `/etc/system` files as you would in standard Sybase ASE installation procedures. For instructions on how to install the Sybase ASE software, refer to the Sybase installation and configuration guides.

Note – For every Sybase server, enter the hostname that is associated with a network resource when asked to specify the hostname.

Where to Go From Here

After you install the Sybase ASE software, go to "How to Configure Sybase ASE Database Access With Solstice DiskSuite/Solaris Volume Manager" on page 17 if you use the Solstice DiskSuite/Solaris Volume Manager. Go to "How to Configure Sybase ASE Database Access With VERITAS Volume Manager" on page 18 if you use the VERITAS Volume Manager (VxVM).

▼ How to Verify the Sybase ASE Installation

Perform the following steps to verify the Sybase ASE software installation.

1. **Verify that the** *sybase* **user and the** *dba* **group own the `$SYBASE_HOME` directory and `$SYBASE_HOME` children directories.**

2. **Run the `scstat(1M)` command to verify that the Sun Cluster software functions correctly.**

Creating the Sybase ASE Database Environment

The procedures in this section enable you to complete the following tasks.

- Configure Sybase ASE database access with Solstice DiskSuite/Solaris Volume Manager or VxVM.
- Create the Sybase ASE database environment.

▼ How to Configure Sybase ASE Database Access With Solstice DiskSuite/Solaris Volume Manager

If you use the Solstice DiskSuite/Solaris Volume Manager, perform the following steps to configure Sybase ASE database access with the Solstice DiskSuite/Solaris Volume Manager.

1. **Configure the disk devices for the Solstice DiskSuite/Solaris Volume Manager software to use.**

   See the *Sun Cluster 3.1 Software Installation Guide* for information on how to configure Solstice DiskSuite/Solaris Volume Manager.

2. **If you use raw devices to contain the databases, run the following commands to change each raw-mirrored metadevice's owner, group, and mode.**

   If you do not use raw devices, do not perform this step.

   a. **If you create raw devices, run the following commands for each device on** *each node* **that can master the Sybase ASE resource group.**

      ```
      # chown sybase /dev/md/metaset/rdsk/dn
      # chgrp dba /dev/md/metaset/rdsk/dn
      ```

```
# chmod 600 /dev/md/metaset/rdsk/dn
```

metaset   Specifies the name of the diskset.

/rdsk/d*n*   Specifies the name of the raw disk device within the *metaset* diskset.

b. Verify that the changes are effective.

```
# ls -lL /dev/md/metaset/rdsk/dn
```

▼ How to Configure Sybase ASE Database Access With VERITAS Volume Manager

If you use VxVM software, perform the following steps to configure Sybase ASE database access with the VxVM software.

1. Configure the disk devices for the VxVM software to use.

See the *Sun Cluster 3.1 Software Installation Guide* for information on how to configure VERITAS Volume Manager.

2. If you use raw devices to contain the databases, run the following commands on the current disk-group primary to change each device's owner, group, and mode.

If you do not use raw devices, do not perform this step.

a. If you create raw devices, run the following command for each raw device.

```
# vxedit -g diskgroup set user=sybase group=dba mode=0600 volume
```

-g *resource-group*   Specifies the name of the resource group. This name can be your choice but must be unique for resource groups within the cluster.

-h *nodelist*   Specifies an optional comma-separated list of physical node names or IDs that identify potential masters. The order here determines the order in which the nodes are considered as primary during failover.

b. Verify that the changes are effective.

```
# ls -lL /dev/vx/rdsk/diskgroup/volume
```

c. Reregister the disk device group with the cluster to keep the VxVM namespace consistent throughout the cluster.

```
# scconf -c -D name=diskgroup
```

▼ How to Create the Sybase ASE Database Environment

Before you perform this procedure, ensure that you have completed the following tasks.

- Establish a highly available IP address and name, that is, a network resource that operates at installation time.

- Locate device paths for all of the Sybase ASE devices—including the master device and system devices—in the highly available local file system or cluster file system. Configure device paths as one of the following file types.

  - regular files
  - raw devices
  - files that the Solstice DiskSuite/Solaris Volume Manager software or the VxVM software manage

- Locate the Sybase ASE server logs in either the cluster file system or the local file system.

- The Sybase ASE 12.0 environment consists of the data server, backup server, monitor server, text server, and XP server. The data server is the only server that you must configure—you can choose whether to configure all of the other servers.

- The entire cluster must contain only one copy of the `interfaces` file. The `$SYBASE` directory contains the `interfaces` file. If you plan to maintain per-node file copies, ensure the file contents are identical.

  All of the clients that connect to Sybase ASE servers connect with Sybase OpenClient libraries and utilities. When you configure the Sybase ASE software, in the `interfaces` file, enter information about the network resource and various ports. All of the clients use this connection information to connect to the Sybase ASE servers.

Perform the following steps to create the Sybase ASE database environment.

1. **Run the GUI-based utility `srvbuild` to create the Sybase ASE database.**

   The `$SYBASE/ASE_12-0/bin` directory contains this utility. See the Sybase ASE document entitled "Installing Sybase Adaptive Server Enterprise on Sun Solaris 2.x (SPARC)."

2. **To verify successful database installation, ensure that all of the servers start correctly.**

   Run the ps(1) command to verify the operation of all of the servers. Sybase ASE server logs indicate any errors that have occurred.

3. **Set the password for the Sybase ASE system administrator account.**

   See the *Sybase Adaptive Server Enterprise System Administration Guide* for details on changing the sa login password.

4. **Create a new Sybase ASE account for fault monitoring.**

   This account enables the fault monitor to perform the following tasks.

   - Support queries to system tables.
   - Create and update user tables.

   Note – Do not use the sa account for these purposes.

   The following example shows how to create a new Sybase ASE account for fault monitoring.

   ```
   # isql -Usa -Psybase -Sasedb
   1> use master
   2> go
   1> create database sc3xdb
   2>go
   1> sp_addlogin dbmon, dbmonp, sc3xdb
   2> go
   1> use sc3xdb
   2> go
   1> sp_changedbowner dbmon
   2> go
   1> sp_modifylogin dbmon, defdb, sc3xdb
   2> go
   1> exit
   ```

   See "Sun Cluster HA for Sybase ASE Fault Monitor" on page 31 for more information.

5. **Update the stop file with the sa password.**

   Because the stop file contains the sa password, protect the file with the appropriate permissions, and place the file in a directory that the system administrator chooses. Enable only the *sybase* user to read, write, and execute the stop file.

   See "Important Security Issues" on page 27 for more information about the stop file.

Where to Go From Here

After you create the Sybase ASE database environment, go to "How to Install Sun Cluster HA for Sybase ASE Packages" on page 21.

Installing the Sun Cluster HA for Sybase ASE Package

You can use the `scinstall`(1M) utility to install `SUNWscsyb`, the Sun Cluster HA for Sybase ASE package, on a cluster. Do not use the `-s` option to non-interactive `scinstall` to install all of the data service packages.

If you installed the `SUNWscsyb` data service package as part of your initial Sun Cluster installation, proceed to "Registering and Configuring Sun Cluster HA for Sybase ASE" on page 22. Otherwise, use the following procedure to install the `SUNWscsyb` package.

▼ How to Install Sun Cluster HA for Sybase ASE Packages

You need the Sun Cluster Agents CD-ROM to complete this procedure. Perform this procedure on all of the cluster nodes that run the Sun Cluster HA for Sybase ASE package.

1. **Load the Sun Cluster Agents CD-ROM into the CD-ROM drive.**

2. **Run the `scinstall` utility with no options.**
   This step starts the `scinstall` utility in interactive mode.

3. **Choose 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 Cluster Agents CD-ROM.**
   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

When you finish the Sun Cluster HA for Sybase ASE package installation, go to "How to Register and Configure Sun Cluster HA for Sybase ASE" on page 22.

Registering and Configuring Sun Cluster HA for Sybase ASE

Use the procedures in this section to register and configure the Sun Cluster HA for Sybase ASE data service. Register and configure Sun Cluster HA for Sybase ASE as a failover data service.

▼ How to Register and Configure Sun Cluster HA for Sybase ASE

This procedure describes how to use the `scrgadm`(1M) command to register and configure Sun Cluster HA for Sybase ASE.

This procedure includes creating the `SUNW.HAStoragePlus` resource type. This resource type synchronizes actions between HAStorage and Sun Cluster HA for Sybase ASE and enables you to use a highly available local file system. Sun Cluster HA for Sybase ASE is disk-intensive, and therefore you should configure the `SUNW.HAStoragePlus` resource type.

See the `SUNW.HAStoragePlus`(5) man page and "Relationship Between Resource Groups and Disk Device Groups" in *Sun Cluster 3.1 Data Service Planning and Administration Guide* for more information about the `SUNW.HAStoragePlus` resource type.

Note – Other options also enable you to register and configure the data service. See "Tools for Data Service Resource Administration" in *Sun Cluster 3.1 Data Service Planning and Administration Guide* for details about these options.

To perform this procedure, you must have the following information.

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

- The network resource that clients use to access the data service. You typically configure the IP address when you install the cluster. See the sections in the *Sun Cluster 3.1 Software Installation Guide* on planning the Sun Cluster environment and on how to install the Solaris operating environment for details.

- The path to the Sybase ASE application installation.

Note – Perform the following steps on one cluster member.

---

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

2. **Run the `scrgadm` command to register resource types for Sun Cluster HA for Sybase ASE.**

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

-a	Adds the resource type for the data service.
-t SUNW.sybase	Specifies the resource type name that is predefined for your data service.

3. **Register the SUNW.HAStoragePlus resource type with the cluster.**

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

4. **Create the resource `sybase-hastp-rs` of type SUNW.HAStoragePlus.**

   ```
   # scrgadm -a -j sybase-hastp-rs -g sybase-rg -t SUNW.HAStoragePlus \
   -x GlobalDevicePaths=sybase-device-group1,/dev/global/dsk/d1 \
   -x FilesystemMountPoints=/global/sybase-inst \
   -x AffinityOn=TRUE
   ```

Note – AffinityOn must be set to TRUE and the local file system must reside on global disk groups to be failover.

---

5. **Run the `scrgadm` command to complete the following tasks and bring the resource group sybase-rg online on a cluster node.**

   - Move the resource group into a managed state.
   - Bring the resource group online

   This node will be made the primary for device group `sybase-set1` and raw device `/dev/global/dsk/d1`. Device groups associated with filesystems such as `/global/sybase-inst` will also be made primaries on this node..

   ```
   # scrgadm -Z -g sybase-rg
   ```

6. **Create Sybase ASE application resources in the failover resource group.**

```
# scrgadm -a -j resource -g resource-group \
-t SUNW.sybase \
-x Environment_File=environment-file-path \
-x Adaptive_Server_Name=adaptive-server-name \
-x Backup_Server_Name=backup-server-name \
-x Text_Server_Name=text-server-name \
-x Monitor_Server_Name=monitor-server-name \
-x Adaptive_Server_Log_File=log-file-path \
-x Stop_File=stop-file-path \
-x Connect_string=user/passwd
-y resource_dependencies=storageplus-resource
```

-j *resource*
  Specifies the resource name to add.

-g *resource-group*
  Specifies the resource group name into which the RGM places the resources.

-t SUNW.sybase
  Specifies the resource type to add.

-x Environment_File=*environment-file*
  Sets the name of the environment file.

-x Adaptive_Server_Name=*adaptive-server-name*
  Sets the name of the adaptive server.

-x Backup_Server_Name=*backup-server-name*
  Sets the name of the backup server.

-x Text_Server_Name=*text-server-name*
  Sets the name of the text server.

-x Monitor_Server_Name=*monitor-server-name*
  Sets the name of the monitor server.

-x Adaptive_Server_Log_File=*log-file-path*
  Sets the path to the log file for the adaptive server.

-x Stop_File=*stop-file-path*
  Sets the path to the stop file.

-x Connect_string=*user/passwd*
  Specifies the user name and password that the fault monitor uses to connect to the database.

You do not have to specify extension properties that have default values. See "Configuring Sun Cluster HA for Sybase ASE Extension Properties" on page 27 for more information.

7. Enable the resource and fault monitoring.

> **Note –** Sybase start logs print to the console when the Sybase servers start. If you do not want these messages to print to the console, update the appropriate RUN files to redirect these messages to another file.

```
# scswitch -Z -g resource-group
```

Where to Go From Here

After you register and configure Sun Cluster HA for Sybase ASE, go to "How to Verify the Sun Cluster HA for Sybase ASE Installation " on page 25.

Verifying the Sun Cluster HA for Sybase ASE Installation

Perform the following verification tests to ensure that you have correctly installed and configured Sun Cluster HA for Sybase ASE.

These sanity checks ensure that all of the nodes that run Sun Cluster HA for Sybase ASE can start the Sybase ASE data server. These checks also ensure that other nodes in the configuration can access the Sybase ASE data server. Perform these sanity checks to isolate any problems with starting the Sybase ASE software from Sun Cluster HA for Sybase ASE.

▼ How to Verify the Sun Cluster HA for Sybase ASE Installation

1. **Log in to the node that masters the Sybase ASE resource group.**

2. **Set the Sybase ASE environment variables.**

   The environment variables are the variables that you specify with the Environment_file extension property. See Table 1–2 for information on setting these environment variables.

3. **Verify that the Sun Cluster HA for Sybase ASE resource is online.**

   ```
   # scstat -g
   ```

4. **Inspect the Sybase ASE logs to determine the cause of any errors that have occurred.**

5. **Confirm that you can connect to the data server and execute the following test command.**

   ```
   # isql -S adaptive-server -U sa -P password

   isql> sp_help
   isql> go
   isql> quit
   ```

6. **Kill the process for the Sybase ASE data server.**

   The Sun Cluster software restarts the process.

7. **Switch the resource group that contains the Sybase ASE resource to another cluster member.**

   ```
   # scswitch -z -g resource-group -h node
   ```

8. **Log in to the node that now contains the resource group.**

9. **Repeat Step 3 and Step 5.**

   Note – Sybase ASE client connections cannot survive a Sun Cluster HA for Sybase ASE switchover. If a switchover occurs, the existing client connections to Sybase ASE terminate, and clients must reestablish their connections. After a switchover, the time that is required to replay the Sybase ASE transaction log determines Sun Cluster HA for Sybase ASE recovery time.

Understanding Sun Cluster HA for Sybase ASE Logging and Security Issues

The following sections contain information about Sun Cluster HA for Sybase ASE logging and security issues.

Sun Cluster HA for Sybase ASE Logging

Sun Cluster HA for Sybase ASE logs messages to the file `message_log` in the `/opt/SUNWscsyb/l` directory. Although this file cannot exceed 512 Kbytes, Sun Cluster HA for Sybase ASE does not delete old log files. The number of log files, therefore, can grow to a large number.

Sun Cluster HA for Sybase ASE writes all of the error messages in the `syslog` file. Sun Cluster HA for Sybase ASE also logs fault monitor history to the file `restart_history` in the `log` directory. These files can also grow to a large number.

As part of your regular file maintenance, check the following log files and remove files that you no longer need.

- `syslog`
- `message_log`
- `restart_history`

Important Security Issues

Sun Cluster HA for Sybase ASE requires that you embed the system administrator's password in a stop file. The `/opt/SUNWscsyb/bin` directory contains the template for the stop file, `sybase_stop_servers`. Sun Cluster HA for Sybase ASE uses this file to log in to the Sybase ASE environment and to stop the Sybase ASE servers. Enable the *sybase* user to execute the stop file, but protect the file from general access. Give read, write, and execute privileges to only the following users.

- *sybase* user
- *sybase* group

Configuring Sun Cluster HA for Sybase ASE Extension Properties

This section describes how to configure Sun Cluster HA for Sybase ASE extension properties. Typically, you use the command line `scrgadm -x` *parameter=value* to configure extension properties when you create the Sybase ASE resources. You can also use the procedures that "Administering Data Service Resources" in *Sun Cluster 3.1 Data Service Planning and Administration Guide* describes to configure them later.

See the `r_properties`(5) and the `rg_properties`(5) man pages for details on all of the Sun Cluster extension properties.

Table 1–2 describes the extension properties that you can set for the Sybase ASE server resource. You can update some extension properties dynamically. You can update others, however, only when you create or disable a resource. The Tunable entries indicate when you can update each property.

TABLE 1–2 Sun Cluster HA for Sybase ASE Extension Properties

Name/Data Type	Description
Environment_File	File that contains all of the Sybase ASE environment variables. Sun Cluster HA for Sybase ASE requires that you define the variables SYBASE, SYBASE_ASE, and SYBASE_OCS. Other variables that you define are passed as environment variables to the Sybase server.
	The definition of each environment variable defined must follow the format
	VARIABLE_NAME=VARIABLE_VALUE
	Each of these environment variables must also be specified, one per line in the Environment_File.
	Typically, users use the SYBASE.sh environment file created by the Sybase installation.
	Note – The value of this property is independent of the shell being used by the *sybase* user. In other words, the *sybase* user can have csh as its default shell.
	Default: None
	Range: Minimum=1
	Tunable: When disabled
Adaptive_Server_Name	The name of the data server. Sun Cluster HA for Sybase ASE uses this property to locate the RUN server in the $SYBASE/$ASE/install directory.
	Default: None
	Range: Minimum=1
	Tunable: When disabled

TABLE 1–2 Sun Cluster HA for Sybase ASE Extension Properties   *(Continued)*

Name/Data Type	Description
Backup_Server_Name	The name of the backup server. Sun Cluster HA for Sybase ASE uses this property to locate the RUN server in the $SYBASE/$ASE/install directory. If you do not set this property, Sun Cluster HA for Sybase ASE will not manage the server. **Default:** Null **Range:** None **Tunable:** When disabled
Monitor_Server_Name	The name of the monitor server. Sun Cluster HA for Sybase ASE uses this property to locate the RUN server in the $SYBASE/$ASE/install directory. If you do not set this property, Sun Cluster HA for Sybase ASE will not manage the server. **Default:** Null **Range:** None **Tunable:** When disabled
Text_Server_Name	The name of the text server. The Sun Cluster HA for Sybase ASE data service uses this property to locate the RUN server in the $SYBASE/$ASE/install directory. If you do not set this property, the Sun Cluster HA for Sybase ASE data service will not manage the server. **Default:** Null **Range:** None **Tunable:** When disabled
Adaptive_Server_Log_ File	The path to the log file for the adaptive server. Sun Cluster HA for Sybase ASE continually reads this property for error monitoring. **Default:** None **Range:** Minimum=1 **Tunable:** When disabled

TABLE 1–2 Sun Cluster HA for Sybase ASE Extension Properties   *(Continued)*

Name/Data Type	Description
Stop_File	Sun Cluster HA for Sybase ASE uses this property during server stoppages. This property contains the sa password. Protect this property from general access. **Default:** None **Range:** Minimum=1 **Tunable:** When disabled
Probe_timeout	Time-out value for the fault monitor probe. **Default:** 30 seconds **Range:** 1 – 99999 seconds **Tunable:** Any time
Debug_level	Debug level for writing to the Sun Cluster HA for Sybase ASE log. **Default:** 0 **Range:** 0 – 15 **Tunable:** Any time
Connect_string	String of format *user/password*. Sun Cluster HA for Sybase ASE uses this property for database probes. **Default:** None **Range:** Minimum=1 **Tunable:** When disabled
Connect_cycle	Number of fault monitor probe cycles before Sun Cluster HA for Sybase ASE establishes a new connection. **Default:** 5 **Range:** 1 – 100 **Tunable:** Any time
Wait_for_online	Whether the START method waits for the database to come online before exiting. **Default:** FALSE **Range:** TRUE – FALSE **Tunable:** Any time

Sun Cluster HA for Sybase ASE Fault Monitor

The Sun Cluster HA for Sybase ASE fault monitor queries the Sybase ASE server to determine server health.

Note – The Sun Cluster HA for Sybase ASE fault monitor only monitors the Adaptive server. The fault monitor does not monitor auxiliary servers.

The fault monitor consists of the following processes.

- a main fault monitor process
- a database-client fault probe

The following sections describe the Sun Cluster HA for Sybase ASE fault monitor processes and the extension properties that the fault monitor uses.

Main Fault Monitor Process

The fault monitor process diagnoses errors and checks statistics. The monitor labels an operation successful if the following conditions occur.

- The database is online.
- The activity check returns no errors.
- The test transaction returns no errors.

If an operation fails, the main process checks the action table for an action to perform and then performs the predetermined action. If an operation fails, the main process can perform the following actions, which execute external programs as separate processes in the background.

1. Restarts the resource on the current node.

2. Restarts the resource group on the current node.

3. Fails over the resource group to the next node on the resource group's nodelist.

The server fault monitor also scans the `Adaptive_Server_Log` file and acts to correct any errors that the scan identifies.

Database-Client Fault Probe

The database-client fault probe performs activity checks and test transactions. The extension property `Connect_string` specifies an account that performs all of the database operations. The extension property `Probe_timeout` sets the time-out value that the probe uses to determine the time that has elapsed in a successful database probe.

Extension Properties

The fault monitor uses the following extension properties.

- `Thorough_probe_interval`
- `Retry_count`
- `Retry_interval`
- `Probe_timeout`
- `Connect_string`
- `Connect_cycle`
- `Adaptive_Server_Log`

See "Configuring Sun Cluster HA for Sybase ASE Extension Properties" on page 27 for more information about these extension properties.

Index

V

verifying
    Sun Cluster HA for Sybase ASE, 25
    Sybase ASE installation, 17