



Technical Report

Integrating Sybase ASE Cluster Edition with NetApp Storage System

Suresh Vundru, NetApp
July 2010 | TR-3857

TABLE OF CONTENTS

1	INTRODUCTION	3
1.1	SCOPE	3
2	ASSUMPTIONS	3
3	INFRASTRUCTURE REQUIREMENTS	3
3.1	CLUSTER NODES	4
3.2	STORAGE INFRASTRUCTURE	4
3.3	SOFTWARE REQUIREMENTS	4
4	INSTALLATION PREREQUISITES	5
4.1	PREINSTALLATION TASKS FOR NETAPP STORAGE	5
4.2	PREINSTALLATION TASKS FOR THE SERVER INSTALLATION	6
5	INSTALLATION PROCEDURE	8
5.1	INSTALLING THE SERVER	8
6	CONFIGURING AND STARTING THE CLUSTER SERVER	10
6.1	BEFORE YOU BEGIN	10
6.2	SETTING THE ENVIRONMENT AND STARTING THE UNIFIED AGENTS	10
7	APPENDIXES	14
7.1	APPENDIX A: SAMPLE CLUSTER CONFIGURATION PARAMETERS AND VALUES	14
7.2	APPENDIX B: SAMPLE TEST CASES TO VERIFY CLUSTER	16
7.3	APPENDIX C: GLOSSARY	24
7	ACKNOWLEDGEMENTS	24

1 INTRODUCTION

This technical report describes the installation and configuration of Sybase® Adaptive Server Enterprise (ASE) Cluster Edition on an Oracle® Solaris SPARC platform with NetApp® storage. This is a certified configuration, so the components presented in this paper must be used in the same combination to gain support from all parties involved. The only exception to this is the application of certain patches, as required by the vendors in this configuration.

1.1 SCOPE

This document describes the following procedures:

- Installing Sybase Adaptive Server Enterprise Cluster Edition with Solaris SPARC on a NetApp storage system
- Configuring and starting the Sybase Cluster Server

For this document, a representative configuration was selected to prove the system from a functionality point of view. Although it is generally desirable to architect a specialized configuration for a particular application, performance specifications and tuning are beyond the scope of this document.

2 ASSUMPTIONS

This document assumes that you are familiar with the following:

- Sybase Adaptive Server
- Operation of NetApp storage systems
- Solaris operating environment

All examples in this technical report were tested using Sybase ASE Cluster Edition 15.0.3 running on a Solaris 10 SPARC platform. The examples in this document may require modification to run in your configuration. This document also assumes that you have access to the Sybase documentation for your particular operating system and that you have followed the documented steps for installation and configuration. If Sybase documentation and this technical report conflict, you should assume that the Sybase documentation is correct. For more information about the administration of the individual products, refer to the related product documentation.

This document also assumes that there are no other Sybase installations on the cluster nodes and that no Sybase processes are running at the time of installation.

3 INFRASTRUCTURE REQUIREMENTS

To set up Sybase Cluster nodes with NetApp storage systems, you must have:

- A Sybase UNIX® server machine running a Sybase-supported operating system (OS); in this case, a Solaris 10 OS SPARC platform
- NetApp storage systems
- Supported network configuration for the previous items
- One NFS mount point with Gigabit Ethernet network connection to install Sybase ASE CE binaries
- For Fibre Channel Protocol, an FCP attachment kit is required from NetApp
- For FC SAN configurations, NetApp highly recommends using the latest NetApp SAN host attach kit for Solaris. The kit comes with the Fibre Channel HBA, drivers, firmware, utilities, and documentation. For installation and configuration, see the documentation that is shipped with the attach kit.

The infrastructure requirements in this section cover a minimal two-node configuration. Although more than two nodes can be used, and for production systems that have redundant interconnects, NetApp recommends network connections for complete redundancy.




Figure 1) System Architecture.

3.1 CLUSTER NODES

- Sun Microsystems SPARC architecture–based servers
- Gigabit Ethernet NICs (for cluster interconnect)
- Gigabit Ethernet NICs (to access NFS)

3.2 STORAGE INFRASTRUCTURE

- Two NetApp storage systems (our configuration uses FAS3070 with Data ONTAP[®] 7.3.1)
- Two or more disk shelves based on disk space requirements
- Two FC switches for redundant connection between storage and cluster nodes
- One gigabit NIC in the storage system

3.3 SOFTWARE REQUIREMENTS

For both the nodes in the participating cluster, unless specified otherwise:

- Solaris 10 (SPARC 64-bit) with the latest patches (our configuration used Solaris 10 5/09 s10s_u7wos_08 SPARC)
- Sybase Adaptive Server Enterprise Cluster Edition 15.0.3 for Solaris (SPARC 64-bit) with latest EBF
- NetApp storage system running Data ONTAP 7.3.x

4 INSTALLATION PREREQUISITES

4.1 PREINSTALLATION TASKS FOR NETAPP STORAGE

CREATING NFS VOLUME

1. Create a new volume on NetApp storage to install Sybase binaries.

For example, to create the volume `sybhome` in `aggr1` aggregate with 50GB size:

```
ControllerA> vol create sybhome aggr1 50g
```

2. Export the `sybhome` volume with `rw` and `root` permission to cluster nodes only using `exportfs`:

```
ControllerA > exportfs -p rw=v440-1:v440-2,root=v440-1:v440-2,anon=0  
/vol/sybhome
```

3. Export the new volume:

```
ControllerA > exportfs -a
```

CREATING FCP LUNS

To create new FCP LUNs on NetApp storage to configure as Sybase devices, follow these steps:

1. Before creating the FCP LUNS, verify that the FCP license installed.
2. Check the status of FCP service on the storage controller:

```
ControllerA> fcp status
```

If FCP service is not started, start it:

```
ControllerA > fcp start
```

3. Create the FCP initiator group:

- i. Check the FCP initiator port names on `rac1` and `rac2`:

```
ControllerA > fcp show initiator
```

The initiator port names should correspond to the Fibre Channel HBA in the cluster nodes.

- ii. Create the initiator group and add initiator port names to it.

For example:

```
ControllerA> igroup create -f -t solaris "v440-all" 1000000c964b256  
1000000c964aaa4 1000000c9583753 1000000c958381f
```

4. Create the LUNS (disks):

The examples in this guide use different LUNs to configure quorum, master, system procedures, and system database devices.

For example:

```
ControllerA> vol create cluster aggr1 10g
```

```
ControllerA> vol create clusterdb aggr1 20g
```

```
lun create -s 2g -t solaris /vol/cluster/quorum.lun
```

```
lun create -s 2g -t solaris /vol/ cluster/master.lun
```

```
lun create -s 2g -t solaris /vol/cluster/spdb.lun
```

```
lun create -s 2g -t solaris /vol/ cluster/ssdb.lun
```

```
lun create -s 10g -t solaris /vol/ clusterdb/testdb.lun
```

5. Map the created LUNs to the `v440-all` initiator group in the example:

```
lun map [ -f ] <lun_path> <initiator_group> [ <lun_id> ]
```

For example:

```
ControllerA> lun map /vol/ clusterdb/testdb.lun v440-all 0
```

6. Check that the LUNS are mapped to the v440-all initiator group:

```
ControllerA>lun show -m
```

4.2 PREINSTALLATION TASKS FOR THE SERVER INSTALLATION

PREPARING DEVICES FOR INSTALLATION

1. Create an NFS mount point (that is, \$SYBASE) for installing the Sybase ASE CE software.

On both cluster nodes, enter:

```
# mkdir /sybhome
```

```
# mount -f nfs -o rw,bg,hard ControllerA:/vol/sybhome /sybhome
```

```
# chown sybase /sybhome
```

Create an entry in /etc/vfstab to make the mount point persistent up on rebooting server

2. Raw devices to be used as Sybase database devices.

Verify that the LUNs created on the storage system are visible to the cluster nodes and that the Sybase user account has read/write access to the devices.

```
# san lun show all
```

3. Use the qrmutil and io fenceutil commands to check that the devices are I/O fence capable.

PREINSTALLATION PROCEDURE

Before installing the server, follow these steps:

1. Install the appropriate operating system patches on each node.
2. Install the latest NetApp SAN host attach kit for Solaris.
3. Shut down all clusters that use the previous version of the Cluster Edition before beginning the software installation.
4. If you are having trouble starting the installer, make sure that you have the required operating system patches for the Java[®] Runtime Environment version 6. Information about the required operating system patches is available from the Oracle Software Downloads site (www.oracle.com/technology/software/index.html).
5. Create a sybase account on your system to perform all installation tasks.
 - This account should own all devices and files, and if must have permission to read and write all devices the cluster uses.
 - Create all devices with the sybase account.
 - Make sure that all disk devices are accessible from all the nodes in the cluster.
 - Make sure that the account that is used to start the cluster has permissions to read and write to all of the disk devices.
 - Make sure that the cluster and the SCSI generic driver have write permission on some of the /dev/rdisk/c* files that correspond to configured database devices.
6. The SCSI driver expects write access to the /dev/ rdsk/c* files for SCSI-3 PGR commands used in I/O fencing. Use the ls -l commands to verify paths and file permissions. Use the dd utility to verify that the sybase account can read and write to the devices.

7. If you need to change the permissions, use `chmod` or `chown` to correct the write permissions for `/dev/rdisk/c*` files. You can change the write access permissions for these files to root only after rebooting the machine.
8. Log in to the machine as the `sybase` user. Maintain consistent ownership and privileges for all files and directories. A single user—the Sybase system administrator with read, write, and execute permissions—should perform all installation, upgrade, and setup tasks.
9. If you are using a local account that local to just one physical box, then create the `sybase` user on each machine and make sure that it has the same user ID.
10. Decide where the Adaptive Server software will be installed.
 - Make sure there is sufficient disk space available.
 - Make sure that there are no spaces in the path name of the directory.
 - If this is a shared installation, the `$SYBASE` location must be on a shared file system, accessible from all cluster nodes that use the same path (in our configuration, NFS mount point).
11. Verify that the operating system meets the version level, RAM, and network protocol requirements for your platform.
12. Verify that your network software is configured. The Cluster Edition implicitly requires your network to be configured for the nodes that you plan to include in the cluster. If you are having connection problems, or to verify your network configuration, ping the host.
13. Adjust the operating system shared memory parameter for your platform, as described next.

ADJUSTING THE OPERATING SYSTEM SHARED MEMORY PARAMETERS

The following example for Oracle Solaris shows how to adjust shared memory parameters for your platform.

To adjust shared memory segments of the operating system, add the following line to the configuration file `/etc/system`, where `x` is the number of shared memory segments:

```
set shmsys:shminfo_shmseg=x
```

To control system resources, add an entry to the `/etc/project` file such as the following, where `project.max-shm-memory=(privileged,17179869184,deny)` is the name for the project parameter:

```
project-sybase:200: For use by Sybase:sybase:sybase:
```

```
project.max-shm-memory=(privileged,17179869184,deny)
```

- **Privileged:** A threshold value on a resource control that constitutes a point at which local actions can be triggered, or global actions, such as logging in to a machine, can occur. You can abbreviate this to `priv`. The three levels of privilege are:
 - **Basic:** The owner of the calling process can modify
 - **Privileged:** Privileged callers can modify
 - **System:** Fixed for the duration of the operating system instance
- **project.max-shm-memory:** The value of the maximum shared memory. The default value for `project.max-shm-memory` is 25% of the physical memory on the system. The maximum value is `UINT64_MAX`, which works out to 18,446,744,073,709,551,615 bytes. Essentially, this is limited only by the size of physical memory. To set the value, use:
 - **prctl:** To set `project.max-shm-memory` while the system is running
 - **rctladm:** to set the values permanently
- **deny:** Attempts to use more than 16GB are denied.

Adjust shared memory segments if Adaptive Server cannot obtain large memory as a single segment or if backup server stripes fail due to insufficient segments. Depending on the number and types of devices you use for backup (dump) and recovery (load), you may need to adjust the shared memory segment parameter in the operating system configuration file to accommodate concurrent backup server processes. There are six default shared memory segments available for process attachments.

If any reconfiguration through `sp_configure` requires additional memory Adaptive Server allocates shared memory segments after startup. To account for these additional segments, allocate the maximum memory available to the Adaptive Server by using the `allocate max shared memory` configuration parameter. For more information, see the system administration guide.

5 INSTALLATION PROCEDURE

5.1 INSTALLING THE SERVER

If you have a working Adaptive Server on your system, do not install the Cluster Edition in the same `$SYBASE` directory. If you install the Cluster Edition into the same directory as earlier versions of the same components, the Cluster Edition overwrites the older versions.

INSTALLING THE SERVER WITH THE INSTALLER

This section describes how to install the server components by using the GUI mode.

Sybase recommends installing products as `sybase` user; however, root permissions are not required to run the installer. The installer creates the target directory (if necessary) and installs the selected components into that directory. At the end of the installation, you can verify the product installation. You may need to perform additional configuration procedures before you can use some products.

You must source the environment variables needed for Adaptive Server products by running the `SYBASE.csh` or `SYBASE.sh` script file after exiting the installer.

1. Stop and shut down all programs before running the installer.
2. Insert the Adaptive Server CD in the appropriate drive, or download and extract the Adaptive Server install image from the Sybase Product Download Center (SPDC).
3. Launch the installer. For a CD drive enter `/cdrom/setup`.
4. Select a directory for the installation. Click Next to accept the default `/opt/Sybase` (in our configuration, change directory to `/sybhome`).
5. In the Install Type window, choose one of the following three types of installation:
 - Typical
 - Full
 - CustomIf you select Custom, the Products and Features Selection window appears. Select the products to install. Some features depend on other features; therefore the installer prevents you from deselecting some items without first deselecting others.
6. Select one of the following installation options for the Adaptive Server suite:
 - **Licensed Copy of Adaptive Server Suite:** Select this option if you have a licensed copy of Adaptive Server.
 - **Evaluation Edition of Adaptive Server Suite:** Select this option to evaluate Adaptive Server.
7. Select your country name from the drop-down list, read the license terms, and then click I Agree. Click Next.
8. One of the following appears:
 - If you selected licensed installation, the installer prompts the SySAM license. In this panel, you can specify the license file, use an existing license server, or proceed with the installation without license.

- If you selected unlicensed installation, the software will work for 30 days from the original installation date.

The Sybase Software Asset Management Notification screen appears.

9. If you specify a server license, the installer prompts to set up the license server. Enter the license key if you have one. Otherwise click Continue. This prompt does not appear if you have selected the Evaluation Edition.

A summary screen appears.

10. Click Next to install the SySAM server.
11. Select the license type for the Cluster Edition:

- **Unknown:** Without a license
- **CP:** CPU license
- **DT:** Development and test license
- **AC:** OEM Application Deployment CPU license
- **LP:** Application-specific CPU license

This prompt does not appear if you have selected the Evaluation Edition.

12. Configuring your server for e-mail notification enables designated users to receive information about license management events that require attention. Provide the following information:

- SMTP server host name
- SMTP server port number
- E-mail return address
- Recipient e-mail addresses
- Message severity that triggers e-mail messages

The Product Selection Summary window displays the selections that you have made.

13. Verify that you have selected the correct type of installation and that you have enough disk space to complete the process. Click Next.

The Install Progress window shows the progress of the installation. The Install Status window displays the result of the installation process. Click Next.

14. If prompted to have an ASE plug-in remember passwords, click Yes or No. Click Next.

The Configure Server Progress window displays the progress of the server configuration. The Installation Complete window is the final window. If you encounter any errors, see the troubleshooting guide.

15. At the end of a successful installation, click Finish.

INSTALLING FROM THE COMMAND LINE

To launch the installer by using the command line interface, follow these steps:

1. Change the default directory to the CD.
2. Run the setup program with the `-console` command line argument:

```
. /setup -is:javaconsole -console
```

In cases where the installer launches automatically, click Cancel to stop the GUI installation, then launch the setup program from a terminal or console.

The flow of the installation is identical to a regular GUI installation, except that the display is written to a terminal window and responses are entered using the keyboard.

6 CONFIGURING AND STARTING THE CLUSTER SERVER

After installing the software, follow the instructions in this section to configure and start the cluster server.

Sybase recommends that you use the Adaptive Server plug-in or the `sybcluster` utility to create and manage your shared-disk cluster. However, you can also configure and manage the cluster manually.

6.1 BEFORE YOU BEGIN

Before you begin, gather the following information:

- Whether the cluster installation will support a shared (NFS) or a private installation.
- Whether the cluster will use LDAP (recommended for private installations) or interfaces files for server lookup information.
- Whether the cluster will support Java in the database via a third-party JVM.
- Whether the cluster will support a Veritas™ Cluster Server (VCS) for managing application servers in the clustered environment.
- Whether the cluster will use a single or multiple backup servers.
- The location of the `$SYBASE` installation directory. If this is to be a private installation, you must provide a separate `$SYBASE` installation directory for each instance.
- The cluster name.
- The number of instances and the instance names.
- The Domain Name Service for the network on which the cluster will run. If the DNS entries are not entered correctly for each node, the Unified Agent and the `sybcluster` utility will not function properly.
- The number of agents in the cluster. Sybase recommends using one node for each of the instances in the cluster with two or more engines per instance.
- The raw devices to be used by each database device. For the Cluster Edition, you must create all devices, database devices, and quorum devices as raw devices on shared disks.
- Local user temporary databases do not require shared storage and can use local file systems created as private devices.
- The range of port numbers used by each instance to exchange messages with other instances via the private interconnection. Choose unique port numbers that are not in use by other applications.
Note: If your system does not have an existing standard, the Adaptive Server plug-in and `sybcluster` supply default values.
- The query or listening port number for each instance. Choose unique port numbers that are not in use by other applications.
- The IP addresses or network names of the private primary and secondary network cards on each node. Only UDP network protocol is supported at this time.

Note: Even though we compared some of the prerequisites with private installation, installation and configuration of private installation setup are beyond the scope of this document.

6.2 SETTING THE ENVIRONMENT AND STARTING THE UNIFIED AGENTS

The Unified Agent enables distributed management of the cluster. Set the `$SYBASE` environment variables and start a Unified Agent on each node that will host instances in the cluster.

The default port for the Unified Agent is 9999. If you prefer to start the agent on a different port, see step 3.

From the `$SYBASE` directory of each instance, follow these steps:

1. Source the `$SYBASE` environment script.
2. Start the Unified Agent:

```
$SYBASE_UA/bin/uafstartup.sh &
```

The Unified Agent creates `$$SYBASE_UA/nodes/<node_name>/log/agent.log`, and sends agent output to that file.

3. Optionally, to specify the UAF port number, where `port_number` is any available network port, use:

```
$$SYBASE_UA/bin/uafstartup.sh -port port_number
```

CREATING A CLUSTER USING AN ASE PLUG-IN

The Adaptive Server plug-in includes a wizard that steps you through the process of creating and starting a cluster.

For complete instructions on using the Adaptive Server plug-in to configure and manage the cluster, see "Administering Clusters with Adaptive Server plug-in" in the Clusters Users Guide. Adaptive Server plug-in and `sybcluster` provide upgrade support.

1. Register the Adaptive Server plug-in on all cluster nodes:

```
$ / $$SYBASE/UAf-2_5/bin/uafstartup.sh -port <port number>
```

2. Start Sybase Central:

```
$ / $$SYBASE/ASEP/aseplugin
```

3. Click the Adaptive Server Enterprise icon in the left pane.
4. Click the Utilities tab in the right pane.
5. Double-click or right-click Create Cluster and select Open to open the Create Cluster wizard. Alternatively, select Tools>Adaptive Server Enterprise>Create Cluster to start the Create Cluster wizard.
6. Enter the name of the cluster.
7. Select the platform from the drop-down list.
8. Choose whether you are creating a Shared or Private installation.
9. Select the maximum number of instances you plan to use for this cluster. Click Next.

Discovery settings: You can change the settings that the Adaptive Server plug-in uses to find nodes to use for the cluster.

For more information see "Administering Clusters" in the Users Guide to Clusters.

10. Select the hosts that participate in the cluster. You can also click Specify Node to specify a node that does not appear in the list of hosts.
11. Enter the Unified Agent user name and password and click Next.

Quorum device path: Enter the full path to the quorum device (`dev/rdisk/cxtxdxsx`).

Click Next.

12. If you are installing a private installation, skip this step until you specify the instances in the cluster. If this is a shared installation, you are prompted at this time to enter the following information:

- Sybase home: The Adaptive Server release directory (also the value of `$$SYBASE`)
- Adaptive Server Enterprise home: Full path to the ASE-15_0 directory (the value of `$$SYBASE/$$SYBASE_ASE`)
- Interfaces directory: Path to the directory containing the interfaces file
- Cluster configuration file: Full path to the cluster configuration file
- Sybase environment script: Full path to the script that you use to set the environment variables
- Start parameters: List any startup parameters for the cluster
- Trace flags: List any trace flags to be used at startup

13. Enter the master device information:

- Master device path: Full path to the master device (`dev/rdisk/cxtxdxss`)
 - Device size: Size in megabytes (MB), gigabytes (GB), or terabytes (TB)
 - Database size: Size in MB, (GB), or TB
 - Page size: Page size of the instance: 2-, 4-, 8-, or 16K (the default is 2K)
- Click Next.

Note: NetApp recommends using 4K page size.

14. If the device you specify does not include SCSI-3 PGR capability, the Adaptive Server plug-in asks if you want to continue. If you select No, the plug-in does not move to the next screen.

15. Enter the system procedures (`sybsystemprocs`) information:

- **System procedures device:** Full path to the `sybsystemprocs` device (`/dev/rdisk/cxtxdxss`)
- **Device size:** Size in MB, GB, or TB
- **Database size:** Size in MB, GB, or TB

Click Next.

16. If the device you specify does not include SCSI-3 PGR capability, the Adaptive Server plug-in asks if you want to continue. If you select No, the plug-in does not move to the next screen.

Enter the system database (`sybssystemdb`) information:

- System database device: Full path to the `sybssystemdb` device (`/dev/rdisk/cxtxdxss`)
- Device size: Size in MB, GB, or TB
- Database size: Size in MB, GB, or TB

If the device you specify does not include SCSI-3 PGR capability, the Adaptive Server plug-in asks if you want to continue. If you select No, the plug-in does not move to the next screen.

Click Next.

17. Select the checkbox to enter a PCI device. Enter the full path to the PCI device and enter the device and database size.

18. Click Next. For information about PCI device, see [Java in Adaptive Server Enterprise](#).

19. If the device you specify does not include SCSI-3 PGR capability, the Adaptive Server plug-in asks if you want to continue. If you select No, the plug-in does not move to the next screen.

20. Verify the following or change to appropriate values:

- Network Protocol: Select the UDP network protocol.
- Starting Port: Select a starting port number for the cluster's network interconnects. Make sure that this port number is not already in use.
- If you have a secondary network available, select the checkbox to provide port details.

21. Click Validate Port to make sure the ports are not currently in use, and then click Next.

22. Specify the instances for the cluster. Click Add to add instances to the list.

23. In the Define Server Instance dialog box, General tab, enter the following information for each instance you are adding to the cluster:

- Instance name: The name of the instance
- Node: Select the node
- Query port: Select the port number for the query port entry in the interfaces file

24. Select Validate Port to make sure that this port is available.

- Log file: Full path to the error log for this instance
 - a) Enter the following information in the Advanced tab:
 - Primary: Select an address from the dropdown list.

- Secondary: Enabled if you specified that a secondary address is available. Select a host name or protocol address for the secondary connect.
- Start parameters: List any startup parameters for the instance.
- Traceflags: List any trace flags to use for the instance at startup.

b) If it is available, enter the following information in the Private configuration tab:

- Sybase home: the Adaptive Server release directory (also the value of \$SYBASE)
- Adaptive Server Enterprise home: Full path to the ASE-15_0 directory (also the value of \$SYBASE/\$SYBASE_ASE)
- Interfaces directory: Path to the directory that contain the interfaces file
- Sybase environment script: Full path to the script that you use to set the environment variables
- Cluster configuration file: Full path to the cluster configuration file

c) Click Next to accept the cluster instance properties.

25. If there is more than one device on the node that supports I/O fencing, the plug-in asks you to confirm the selection.
26. Specify the local system temporary databases and devices. You must create local system temporary database on a shared disk. The plug-in lists the default values for the local temporary databases. Accept the defaults or select each table cell to change its value.
27. To add more devices, click Create Device and enter the following information:
 - Device name: Name of the device
 - Device path: Full path to the database device
 - Device size: Size in MB, GB, or TB
 Click OK and then click Next.
28. Review the information on the Summary page to be sure that it's correct. To automatically start the cluster after the wizard finishes, select the appropriate box when the cluster is created. If this option is not selected, the cluster is left in a shutdown state after the Create Cluster wizard finishes.
29. Click Next to create this cluster.

The Adaptive Server plug-in displays another window showing the server messages as it creates the cluster.

Note: The Adaptive Server plug-in and `sybcluster` do not create a `run_server` file when they create a cluster. After you have created a cluster by using the Adaptive Server plug-in or `sybcluster`, you must start this cluster and each instance by using the Adaptive Server plug-in or the `sybcluster` utility. You cannot start this cluster from the command line by using `run_server` files.

STARTING THE CLUSTER WHEN THE OPERATING SYSTEM STARTS

It is possible to configure the host system operating system to automatically start the Adaptive Server cluster when the operating system starts. To do this, by start the Unified Agent on the host system and then execute the `sybcluster` commands to start the cluster by using a shell script. The following steps are an example:

1. Start UAF (Unified Agent Framework).
2. Verify that UAF has started successfully.
3. Execute `sybcluster` and pass a command file to `sybcluster` containing commands to start the instances on the system.

The command file is passed using the `sybcluster -i` command line parameter and appears as follows:

```
connect to asecel5
start instance asecel
```

quit

4. The **sybcluster** command line looks something like the following, where `asecel_startup` is the command file shown above:

```
sybcluster -U uafadmin -P -F host1:9999,host2:9999 -i asecel_startup
```

7 APPENDIXES

7.1 APPENDIX A: SAMPLE CLUSTER CONFIGURATION PARAMETERS AND VALUES

The cluster is created with the following settings:

Cluster name: `abccluster`

Maximum number of instances: `4`

Installation mode: `Shared`

Membership Mode: `Native`

Participating nodes:

`v440-3:12999`

`v440-4:12999`

Start-up parameters for all servers:

Start parameters:

Trace flag for all servers:

Trace flags:

Network protocol: `udp`

Quorum device: `/dev/rdisk/c5t60A98000486E56724E4A536C36436C38d0s4`

Master device: `/dev/rdisk/c5t60A98000486E56724E4A536C35696E39d0s4`

Device size: `2048`

Database size: `1024`

Page size: `4`

System procedures device: `/dev/rdisk/c5t60A98000486E56724E4A536C36455051d0s4`

Device size: `1024`

Database size: `600`

System database device: `/dev/rdisk/c5t60A98000486E56724E4A536C36457A39d0s4`

Device size: `1024`

Database size: `600`

Define Server Instance #1

Instance name: abc1

Instance ID: 1

Node: v440-3

Agent port: 12999

Query port: 5021

Primary network address: v440-3.private

Secondary network address: v440-3

Sybase home: /sybase

ASE home: /sybase/ASE-15_0

Interfaces file directory: /sybase

Sybase environment script: /sybase/SYBASE.sh

Configuration file: /sybase/abccluster.cfg

Log file: /sybase/ASE-15_0/install/abccluster_abc1.log

Start parameters:

Trace flags:

Local system temporary database:

Database name: lstdb_1

Database size: 30

Device name: master

Define Server Instance #2

Instance name: abc2

Instance ID: 2

Node: v440-4

Agent port: 12999

Query port: 5022

Primary network address: v440-4.private

Secondary network address: v440-4

Sybase home: /sybase

ASE home: /sybase/ASE-15_0

Interfaces file directory: /sybase
Sybase environment script: /sybase/SYBASE.sh
Configuration file: /sybase/abccluster.cfg
Log file: /sybase/ASE-15_0/install/abccluster_abc2.log

Start parameters:

Trace flags:

Local system temporary database:

Database name: lstdb_2

Database size: 30

Device name: master

7.2 APPENDIX B: SAMPLE TEST CASES TO VERIFY CLUSTER

TESTING DIRECT I/O

```
[@v440-3:/sybase]isql -Sabcccluster -Usa -P
```

```
1> select @@instancename
```

```
2> go
```

```
-----  
abc1
```

```
(1 row affected)
```

```
1> sp_helpdevice
```

```
2> go
```

```
device_name physical_name  
description  
status cntrltype vdevno vpn_low vpn_high  
-----  
master /dev/rdisk/c5t60A98000486E56724E4A536C35696E39d0s4  
raw device, special, dsync on, directio off, default disk, physical dis  
k, 2048.00 MB, Free: 946.00 MB  
3 0 0 0 1048575  
sysprocsdev /dev/rdisk/c5t60A98000486E56724E4A536C36455051d0s4  
raw device, special, dsync off, directio off, physical disk, 1024.00 MB  
, Free: 424.00 MB  
2 0 1 0 524287  
systemdbdev /dev/rdisk/c5t60A98000486E56724E4A536C36457A39d0s4  
raw device, special, dsync off, directio off, physical disk, 1024.00 MB  
, Free: 424.00 MB
```



```

                2          0          2          0    524287
tapedump1  /dev/rmt4
            unknown device type, disk, dump device
                16          2          0          0    20000
tapedump2  /dev/rst0
            unknown device type, tape,          625 MB, dump device
                16          3          0          0    20000

```

(5 rows affected)

(return status = 0)

1> use master

2> go

1> disk init

2> name = "abc_db1",

3> physname = "/dev/rdisk/c5t60A98000486E56724E4A53706E2F5A50d0s4",

4> size = '2g', dsync = false

5> go

1> sp_helpdevice abc_db1

2> go

```

device_name physical_name
            description
            status cntrltype vdevno vpn_low vpn_high
-----
abc_db1      /dev/rdisk/c5t60A98000486E56724E4A53706E2F5A50d0s4
            raw device, special, dsync off, directio off, physical disk, 2048.00 MB
            , Free: 2048.00 MB
                2          0          3          0    1048575

```

(1 row affected)

(return status = 0)

1> create database mydb1 on abc_db1 = '1g'

2> go

```

CREATE DATABASE: allocating 262144 logical pages (1024.0 megabytes) on disk
'abc_db1' (262144 logical pages requested).
Processed 103 allocation unit(s) out of 1024 units (allocation page 26112). 10%
completed.
Processed 205 allocation unit(s) out of 1024 units (allocation page 52224). 20%
completed.

```

```
Processed 308 allocation unit(s) out of 1024 units (allocation page 78592). 30%
completed.
Processed 410 allocation unit(s) out of 1024 units (allocation page 104704). 40%
completed.
Processed 512 allocation unit(s) out of 1024 units (allocation page 130816). 50%
completed.
Processed 615 allocation unit(s) out of 1024 units (allocation page 157184). 60%
completed.
Processed 717 allocation unit(s) out of 1024 units (allocation page 183296). 70%
completed.
Processed 820 allocation unit(s) out of 1024 units (allocation page 209664). 80%
completed.
Processed 922 allocation unit(s) out of 1024 units (allocation page 235776). 90%
completed.
Processed 1024 allocation unit(s) out of 1024 units (allocation page 261888).
100% completed.
Database 'mydb1' is now online.
```

```
1> use mydb1
2> go
1> create table publ (number int, name varchar(100))
2> go

1> insert into publ (number, name) values(123, "Teating Sybase CE with NetApp Storage")
2> go 10000
(1 row affected)
10000 xacts:
----->> time taken 12sec

1> sp_helpdevice abc_db1
2> go

device_name physical_name
            description
            status cntrltype vdevno vpn_low vpn_high
-----
abc_db1      /dev/rdisk/c5t60A98000486E56724E4A53706E2F5A50d0s4
            raw device, special, dsync off, directio off, physical disk, 2048.00 MB
            , Free: 1024.00 MB
```

```
2 0 3 0 1048575
```

```
(1 row affected)
```

```
dbname size          allocated          vstart lstart
-----
mydb1   1024.00 MB Nov 7 2009 5:15AM 0 0
```

```
(1 row affected)
```

```
(return status = 0)
```

```
1> sp_deviceattr abc_db1, directio, true
```

```
2> go
```

```
'directio' attribute of device 'abc_db1' turned 'on '. Restart Adaptive Server
for the change to take effect.
```

```
(return status = 0)
```

```
1> sp_helpdevice abc_db1
```

```
2> go
```

```
device_name physical_name
description
status cntrltype vdevno vpn_low vpn_high
-----
abc_db1      /dev/rdisk/c5t60A98000486E56724E4A53706E2F5A50d0s4
raw device, special, dsync off, directio on, physical disk, 2048.00 MB,
Free: 1024.00 MB
2 0 3 0 1048575
```

```
(1 row affected)
```

```
dbname size          allocated          vstart lstart
-----
mydb1   1024.00 MB Nov 7 2009 5:15AM 0 0
```

```
(1 row affected)
```

```
(return status = 0)
```

```
----->>> Restarted Adaptive Server for the change to take effect
```

```
[@v440-3:/sybase]isql -Sabcccluster -Usa -P
```

```
1> select @@instancename
```

```
2> go
```

abc1

(1 row affected)

1> use mydb1

2> go

1> sp_helpdevice abc_db1

2> go

device_name	physical_name	description	status	cntrltype	vdevno	vpn_low	vpn_high
-------------	---------------	-------------	--------	-----------	--------	---------	----------

abc_db1	/dev/rdisk/c5t60A98000486E56724E4A53706E2F5A50d0s4	raw device, special, dsync off, directio on, physical disk, 2048.00 MB,					
		Free: 1024.00 MB					
			2	0	3	0	1048575

(1 row affected)

dbname	size	allocated	vstart	lstart
mydb1	1024.00 MB	Nov 7 2009 5:15AM	0	0

(1 row affected)

(return status = 0)

1> insert into publ (number, name) values(123, "Teating Sybase CE with NetApp Storage")

2> go 10000

(1 row affected)

10000 xacts:

----->> time taken 10sec

TESTING DISK INIT

```
[@v440-3:/sybase]isql -Sabcccluster -Usa -P
1> select @@instancename
2> go
-----
      abc1
(1 row affected)
1> use master
2> go
1> disk init
2> name = "abc_db2",
3> physname = "/dev/rdisk/c5t60A98000486E56724E4A53706E2F5A50d0s5",
4> size = '2g'
5> go
```

TESTING CLUSTER FAILOVER

```
[@v440-3:/sybase]isql -Sabcccluster -Usa -P
1> select @@instancename
2> go
-----
      abc1
(1 row affected)
1> use master
2> go
1> create database mydb2 on abc_db2 = '1g'
2> go
CREATE DATABASE: allocating 262144 logical pages (1024.0 megabytes) on disk
'abc_db2' (262144 logical pages requested).
Processed 103 allocation unit(s) out of 1024 units (allocation page 26112). 10%
completed.
Processed 205 allocation unit(s) out of 1024 units (allocation page 52224). 20%
completed.
Processed 308 allocation unit(s) out of 1024 units (allocation page 78592). 30%
completed.
Processed 410 allocation unit(s) out of 1024 units (allocation page 104704). 40%
completed.
Processed 512 allocation unit(s) out of 1024 units (allocation page 130816). 50%
completed.
Processed 615 allocation unit(s) out of 1024 units (allocation page 157184). 60%
```

```

completed.
Processed 717 allocation unit(s) out of 1024 units (allocation page 183296). 70%
completed.
Processed 820 allocation unit(s) out of 1024 units (allocation page 209664). 80%
completed.
Processed 922 allocation unit(s) out of 1024 units (allocation page 235776). 90%
completed.
Processed 1024 allocation unit(s) out of 1024 units (allocation page 261888).
100% completed.
Database 'mydb2' is now online.
1> use mydb2
2> go
1> sp_helpdevice abc_db2
2> go
device_name physical_name
description
status cntrltype vdevno vpn_low vpn_high
-----
-----
abc_db2      /dev/rdisk/c5t60A98000486E56724E4A53706E2F5A50d0s5
raw device, special, dsync off, directio off, physical disk, 2048.00 MB
, Free: 1024.00 MB
          2          0          4          0 1048575

(1 row affected)
dbname size          allocated          vstart lstart
-----
mydb2    1024.00 MB Nov  7 2009  6:45AM      0      0

(1 row affected)
(return status = 0)
1> create table pub2 (num int, name varchar(100))
2> go
1> insert into pub2 (num, name) values(555, "To test failover of the Cluster")
2> go 100
(1 row affected)
100 xacts:

1> shutdown abcl

```

```

2> go
CT-LIBRARY error:
      ct_results(): network packet layer: internal net library error: Net-Library
operation terminated due to disconnect
[@v440-3:/sybase]

----->>> node2

[@v440-4:/sybase/UAF-2_5/bin]isql -Sabccluster -Usa -P
1> select @@instancename
2> go

-----
abc2
(1 row affected)
1> use mydb2
2> go
1> sp_tables pub2
2> go

table_qualifier          table_owner
      table_name
      table_type  remarks
-----
mydb2                dbo
      pub2
      TABLE          NULL

(1 row affected)
(return status = 0)

1> select count(*) from pub2
2> go

-----
      100
(1 row affected)

```

TESING I/O FENCING

```
1> sp_configure 'enable i/o fencing'
```

```
2> go
```

```

Parameter Name          Default      Memory Used Config Value
      Run Value      Unit          Type
      Instance Name
-----
enable i/o fencing          0          0          0
              0 switch          static cluster
              NULL

```

```
(1 row affected)
```

```
(return status = 0)
```

7.3 APPENDIX C: GLOSSARY

Table 1) Glossary.

Term	Description
cluster	A collection of homogeneous nodes in a network that operate as a single system. Each node has its own CPU and memory. All nodes communicate with each other through private and high-speed communication pathways.
clusterware	Sybase software included in Adaptive Server that enables the shared-disk cluster.
failover	The ability to switch automatically to another instance upon the failure or abnormal termination of a previously active node.
instance	An Adaptive Server that participates in a shared-disk cluster.
instance number	A number that uniquely identifies a named instance in the Adaptive Server shared-disk cluster.
local system temporary database	Space for temporary tables and worktables. Each instance in the cluster has a local system temporary database that it alone can access.
node	A machine (hardware) that hosts an Adaptive Server instance.
quorum device	This device provides important information that defines the cluster, including the name of the cluster, the names of the instances in the cluster, the number of nodes, and their names. In addition, the quorum device holds state information about the instances in the cluster and defines cluster membership.

7 ACKNOWLEDGEMENTS

The author would like to thank the following individuals for their contributions to this technical report:

NetApp

Michelle Nguyen, Sridhara Gangoor

Sybase

Amit Satoor, Dennis Yang, Jing He, James Yu Wang

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.



www.netapp.com

© Copyright 2010 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, and Data ONTAP are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Oracle and Java are registered trademarks of Oracle Corporation. UNIX is a registered trademark of The Open Group. Veritas is a trademark of Symantec Corporation. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-3857