

# Sun™ Cluster with NetApp Storage (NFS) for High Availability

Bijaya Kumar Meher, Network Appliance, Inc.

June 2007 | TR-3570

## Executive Summary

In today's enterprise environment, having access to data depends heavily on IT infrastructures that are reliable and resilient enough to adapt to changing conditions without stopping operations. As the demand for high-availability systems increases, Network Appliance, Inc explored the technology necessary to support such infrastructure with its storage product line. NetApp, in collaboration with Sun Microsystems, delivers a solution for Sun Cluster with NetApp storage systems to run today's mission- and business-critical applications. This technical report discusses the procedures to complete the installation of this solution.

## Table of Contents

---

1. Introduction .....	2
2. Assumptions .....	2
3. Support of Fencing and Quorum .....	2
4. Solution and its Components .....	3
5. Servers and Network Appliance Storage in HA environment .....	4
6. Getting the support package .....	5
6.1. <i>Installing the package</i> .....	6
6.2. <i>Removing the package</i> .....	6
6.3. <i>Host Requirement</i> .....	7
6.4. <i>Storage system requirements</i> .....	7
7. Storage System Configuration .....	8
7.1. <i>Connecting with the storage system</i> .....	8
7.2. <i>Register the storage system information</i> .....	9
7.3. <i>Check for iSCSI License</i> .....	10
7.4. <i>Start iSCSI if it is not running</i> .....	10
7.5. <i>Get the iSCSI node name (optional)</i> .....	10
8. Setting up a LUN as a Quorum Device .....	11
8.1. <i>Overview of creating a LUN on a storage system</i> .....	11
8.2. <i>Requirements for the LUN used as a quorum device</i> .....	11
8.3. <i>Setting up a LUN to be a quorum device</i> .....	12
9. Configuring Network appliance storage (nfs) volume for HA File System .....	15
9.1. <i>Creating volumes on Network appliance storage</i> .....	15
9.2. <i>Adding exported directories from NAS device to sun cluster</i> .....	16
10. Configuring Sun cluster with Network appliance storage volume (nfs) for Applications: .....	17
10.1. <i>Pre-Installation Requirements:</i> .....	17
10.2. <i>Creating Application resource group:</i> .....	18
11. References .....	20
12. Acknowledgements .....	20
13. Disclaimer .....	20

## 1. Introduction

This technical report documents installation and configuration of High Availability Sun Cluster version 3.1 and 3.2 software on Solaris™ operating system on Sparc and AMD Opteron™ x64 platforms with NetApp storage in an NFS environment.

The NetApp NFS Client Support Package for Sun Cluster (support package) is a tool that works with NetApp storage systems running Data ONTAP® software that are used in a Sun Cluster environment. In this environment the clustered hosts access shared storage on clustered NetApp storage systems over the Network File System (NFS). With the NetApp support package, you can use the unified storage capability of a NetApp storage system to have an IP-based solution without an external shared SCSI or FCP disk.

## 2. Assumptions

This document assumes that the reader is already familiar with the common administrative tasks related to the Sun Cluster, its concept of quorum and I/O fencing, the NetApp storage systems, SCSI-3, and iSCSI concepts. For more information about the administration of the individual products, please refer to the related product documentation. However, some administrative functions will be described when deemed appropriate for the understanding of the tools.

## 3. Support of Fencing and Quorum

When you install the support package on Sun Cluster hosts running your application, it performs the following two tasks:

- In the event of a host failure, the support package prevents the failed host from writing data to the storage systems. It fences off the failed host from the storage systems.
- The support package also lets the Sun Cluster hosts use a logical unit number (LUN) on the storage systems as a quorum device. Having the quorum capabilities ensures that multiple independent clusters do not exist if a host cluster failure occurs. The Sun Cluster uses SCSI-3 persistent reservation mechanism to maintain membership information about the cluster on the LUN.

## 4. Solution and Its Components

The Sun Cluster system attempts to prevent data corruptions and ensure data integrity by the use of quorum device and I/O fencing. The quorum device is used when a cluster becomes partitioned into separate sets of nodes, to establish which set of nodes constitutes the new cluster. In other words, it ensures multiple independent clusters do not exist in case of a cluster failure. Once the quorum device helped determine which cluster is active, it fences off I/O from any node that does not belong to the cluster. This I/O fencing mechanism ensures that any failed node do not have access to the shared data.

NetApp uses the storage's LUN as a quorum device and iSCSI as the transport between Sun Cluster hosts and the NetApp storage systems.

Hence, the software supplied by NetApp consists of two tools.

1. The first part accomplishes the fencing of the cluster nodes from the data exported by the storage system. This will be shipped in the form of a binary called **NTAPfence**.
2. The second part enables the usage of the filer's LUN as the quorum device. This will be shipped in the form of a binary called **NTAPquorum**.

**NTAPfence** and **NTAPquorum** will be supplied as part of the Solaris **NTAPc1nas** package. This package is available from the NOW™ Website at <http://now.netapp.com>. The package can be installed by using the solaris pkgadd command:

```
# pkgadd -d . NTAPc1nas
```

The binaries NTAPfence and NTAPquorum will be installed in the /usr/sbin directory. The package can be removed by using the solaris pkgrm command:

```
# pkgrm NTAPc1nas
```

## 5. Servers and Network Appliance Storage in HA environment

Figure-1 shows a high availability configuration consisting of two Solaris systems running Sun Cluster connected via NFS through two Gigabit Ethernet switches to a NetApp storage system cluster.

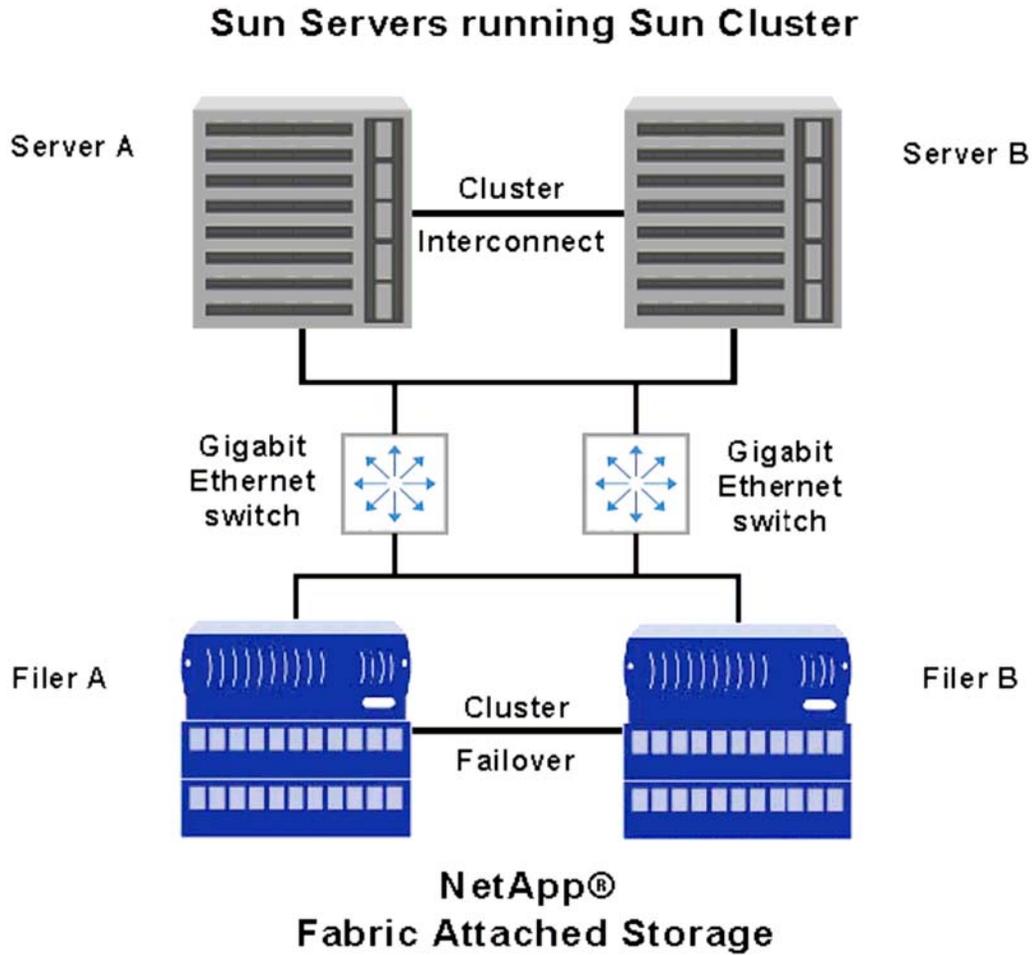


Figure-1

Since this is a high-availability configuration, if one server (for example, Server B) fails, the other server (Server A) runs the support package. This package executes the NTAPfence tool, which rewrites the `/etc/exports` file on Filer A and on Filer B so that Server B cannot write to storage on either storage system.

## 6. Getting the Support Package

To get the support package, you must download the Solaris NTAPcInas package from the NOW site ([now.netapp.com](http://now.netapp.com)). This package contains two binaries:

- NTAPfence
- NTAPquorum

### Downloading the Software

To download the support package software from NOW, complete the following steps

Step	Action
1.	Log in to <a href="http://now.netapp.com">http://now.netapp.com</a> . In the fields on the right side of the page, complete the following steps:  a. Enter your NetApp NOW site user name and password.  b. From the Software download list box, click Download Software.
2.	Go to Package for Sun Cluster and NetApp NFS row of the Software Download table and select Solaris from the Select Platform drop-down list.
3.	Click Go!  <b>Note:</b> Sun Cluster and NetApp NFS agent 1.1 is supported on the sparc platform and NetApp NFS agent 1.2 is supported on x86/x64 platform.
4.	Follow the prompts to reach the Software Download page.
5.	Download the Solaris NTAPcInas package to a local directory on your host machine.

## 6.1. Installing the Package

After you've downloaded the NTAPclnas package, install it by completing the following steps.

**Note:**

You must install this package on each host in the Sun Cluster.

Step	Action
1.	Go to one of the Sun Cluster hosts and enter the following command:  <pre># pkgadd -d NTAPclnas</pre> This installs the binaries NTAPfence and NTAPquorum in the directory /usr/sbin. These are the programs that support package uses.
2.	Repeat Step 1 on each host in the cluster.

## 6.2. Removing the Package

If you want to remove the package, complete the following steps.

Step	Action
1.	Go to one of the Sun Cluster hosts and enter the following command:  <pre># pkgrm NTAPclnas</pre>
2.	Repeat Step 1 on each host in the cluster.

### 6.3. Host Requirement

The following table lists the minimum requirements for using the support package.

Component	Requirement
Operating System	Solaris version 8, 9, and 10 on SPARC® platform. Solaris version 10 update2 on AMD Opteron™ x64 platform.
Sun Cluster software	Sun cluster version 3.1 and 3.2. Check the Software Description page for the support package on the <a href="http://now.netapp.com">http://now.netapp.com</a> site to determine the requirements for the Sun Cluster software.

### 6.4. Storage System Requirements

Your storage system must meet the requirements listed in the Data ONTAP documentation as well as the following requirements. If you have a pair of clustered storage systems, these requirements apply to each storage system.

Component	Minimum Requirement
Operating system	Data ONTAP 7.0.1 or later
Storage appliance setup	With clustered storage systems, make sure you specified the partner IP address in the storage system cluster to be used if a storage system failover occurs.  <b>Note:</b> You specify this address when you run the setup program on the storage system, so you should not need to set it now. See “Confirm the storage system has a partner IP address”. You do not need to specify this information if you have unclustered storage systems.
Licenses	iSCSI license. The LUN used as a quorum device must use iSCSI.  <b>Note:</b> Make sure and check if the iSCSI protocol is running on the storage system that contains the LUN you’re using as a quorum device.

## 7. Storage System Configuration

### 7.1. Connecting with the Storage System

Many of the setup procedures in this section require that you connect with the storage system and log in as root.

There are two primary ways to connect with your storage system:

- Use the command line directly at the console of your storage system.
- Use a command such as telnet to access the storage system.

To use the telnet command to access your storage system, complete the following steps.

Step	Action
1.	<p>Enter the telnet command with either the storage system name or the storage system IP address.</p> <p>If your system configuration does not allow you to log in as root, use the su command to become root once you have accessed the storage system.</p> <p><b>Example:</b> This example opens a Telnet session from the host to the filer via its name, btc-pp-e-filer7:</p> <pre># telnet btc-pp-e-filer7</pre> <p>This is the same storage system, but this time the telnet command uses the system's IP address, 10.73.68.107.</p> <pre># telnet 10.73.68.107</pre>

To set up HTTP on your storage system, complete the following steps:

Step	Action
1.	<p>Set the password for the user sending requests to the storage system (Normally, this is the root user). Use the <code>passwd</code> command to set the password.</p> <p><b>Note:</b> This username/password must be the same one that you provided when you registered the storage system with the Sun Cluster software using the <code>scnas</code> command.</p>
2.	<p>Enable HTTP admin access on the storage system by turning on <code>httpd.admin.access</code> for all the hosts that require admin access.</p> <p><b>Example:</b> The following commands enable HTTP admin access on all hosts.</p> <pre>filer&gt; options httpd.admin. httpd.admin.access all httpd.admin.enable on</pre> <p>For more information on administrative commands, refer to the Data ONTAP documentation.</p>

## 7.2. Register the Storage System Information

Step	Action
1.	<p><b>a.</b> From any cluster node, add the device.</p> <ul style="list-style-type: none"><li>If you are using Sun Cluster 3.2, use the following command: <pre># clnasdevice add -t netapp -p userid=root myfiler</pre><p>Please enter password</p><p>-t netapp      Enter netapp as the type of device you are adding. -p userid=root Enter the HTTP administrator login for the NAS device. myfiler        Enter the name of the NAS device you are adding.</p></li><li>If you are using Sun Cluster 3.1, use the following command: <pre># scnas -a -h myfiler -t netapp -o userid=root</pre><p>Please enter password</p><p>-a              Add the device to cluster configuration. -h myfiler      Enter the name of the NAS device you are adding. -o userid=root   Enter the HTTP administrator login for the NAS device.</p></li></ul> <p><b>b.</b> At the prompt, type the HTTP administrator password.</p> <p><b>c.</b> Confirm that the device has been added to the cluster.</p> <ul style="list-style-type: none"><li>If you are using Sun Cluster 3.2, use the following command: <pre># clnasdevice list</pre><p>For more information about the clnasdevice command, see the clnasdevice (1CL) man page.</p></li><li>If you are using Sun Cluster 3.1, use the following command: <pre># scnas -p</pre></li></ul>

### 7.3. Check for iSCSI License

The LUN designed as a quorum device uses the iSCSI protocol. To determine whether iSCSI is licensed on the storage system containing that LUN, complete the following step.

Step	Action
1.	At the storage system prompt, check the status of iSCSI by entering the following command:  <code># iscsi status</code>  If iSCSI is not licensed on that storage system, a message similar to the following appears:  <code>iscsi: iSCSI is not licensed.</code>

### 7.4. Start iSCSI if It Is Not Running

To start iSCSI on your storage appliance if it is licensed but not running, complete the following step.

Step	Action
1.	Enter the following command:  <code># iscsi start</code>

### 7.5. Get the iSCSI Node Name (Optional)

When you set up the LUN as a quorum device, you can use either the iSCSI node name or the iSCSI host name. If you want to use the node name, use the iSCSI node name on the storage system.

**Example:** This example gets the iSCSI Qualified Name (IQN) of the storage system `btc-ppe-filer7`.

```
btc-ppe-filer7*> iscsi nodename
iSCSI target nodename: iqn.1992-08.com.netapp:sn.33583581
```

You can then set the host name as the suffix for the IQN. The below command is an example of replacing host name, `btc-ppe-filer7`, as the suffix for the IQN of the storage system.

```
btc-ppe-filer7*> iscsi nodename iqn.1992-08.com.netapp: btc-ppe-filer7
btc-ppe-filer7*> iscsi nodename
iSCSI target nodename: iqn.1992-08.com.netapp: btc-ppe-filer7
```

## 8. Setting Up a LUN as a Quorum Device

### 8.1. Overview of Creating a LUN on a Storage System

Step	Action
1.	<p>On the storage system, create an igroup for each storage system in the cluster. The igroup is a collection of the host WWPNs. Use the following command:</p> <pre>igroup create { -f   -i } [ -t &lt;ostype&gt; ] &lt;initiator group&gt; [ &lt;node&gt; ... ]</pre>
2.	<p>Create a LUN using the LUN create command, which has the following format:</p> <pre>lun create -s &lt;size&gt; [-t &lt;type&gt;] [-o noreserve] &lt;lun_path&gt;</pre> <p>If you do not want space reservations on, be sure you include the <code>-o noreserve</code> option on the command line.</p> <p><b>Note</b> You can also create a LUN interactively using the <code>lun setup</code> command.</p>

### 8.2. Requirements for the LUN Used as a Quorum Device

The iSCSI group used for this LUN must be exclusively reserved for this purpose.

**Note:**

Only the iSCSI initiator node names of the cluster nodes can be in the group. This is because if someone else uses the LUN, the quorum information could be destroyed.

### 8.3. Setting up a LUN to be a Quorum Device

Step	Action
1.	Connect to the storage system and login as root.
2.	<p>Display a list of all the LUNs using the <code>lun show -m</code> command.</p> <p><b>Example:</b> The <code>lun show -m</code> command displays the information about the LUNs.</p> <pre> btc-ppe-filer7*&gt; lun show -m LUN path                               Mapped to           LUN ID  Protocol ----- /vol/testflex/testlun0                 testgroup            0       iSCSI /vol/vol0/clustervoll                  clusteragent         0       iSCSI </pre>
3.	If no LUN is listed for your cluster, follow the next step to create a LUN.
4.	<p>To create a LUN for your cluster, use the <code>lun create</code> command. You can also create a LUN interactively using the <code>lun setup</code> command.</p> <p><b>Example:</b> The following command line creates a LUN named <code>clusterA.lun</code> that is 1 GB (-s 1g) in size and has a path of <code>/vol/vol0/</code>.</p> <pre> btc-ppe-filer7*&gt;lun create -s 1g -t solaris /vol/vol0/clusterA.lun </pre>
5.	<p>Confirm that the new LUN exists by entering the <code>lun show</code> command.</p> <p><b>Example:</b> The <code>lun show</code> command displays a list that includes the LUN you created in Step 4.</p> <pre> btc-ppe-filer7*&gt; lun show /vol/vol0/clusterA.lun    1g (1073741824)      (r/w, online) /vol/testflex/testlun0   300.0g (322163441664) (r/w, online, mapped) /vol/vol0/clustervoll    100m (104857600)    (r/w, online, mapped) </pre>
6.	<p>Display a list of all the igroups by entering the <code>igroup show</code> command.</p> <pre> btc-ppe-filer7*&gt; igroup show clusteragent (iSCSI) (ostype: default):   iqn.1986-05.com.sun:cluster-iscsi-quorum.btc-ppe-srv10 (not logged in)   iqn.1986-05.com.sun:cluster-iscsi-quorum.btc-ppe-srv9 (not logged in) testgroup (iSCSI) (ostype: windows):   iqn.1991-05.com.microsoft:btc-ppe-dc.btcppe.netapp.com (logged in on: e0a) </pre>
7.	If your cluster is not listed, go to step 8 and follow the steps to create a new igroup.
8.	<p>Create an igroup for your LUN by using the <code>igroup create <i>clustername</i></code> command.</p> <p><b>Example:</b> The following command line creates an igroup named <code>clusterA</code>.</p> <pre> btc-ppe-filer7*&gt; igroup create -i -t solaris clusterA </pre>

	<p>The OS type can be either “Solaris”, “windows”, “hpux”, “aix”, “linux”, “netware” or “vmware”. The type applies to all initiators within the group and governs the finer details of SCSI protocol interaction with these initiators.</p> <pre> btc-ppe-filer7*&gt; igroup show <b>clusterA (iSCSI) (ostype: solaris):</b>   clusteragent (iSCSI) (ostype: default):     iqn.1986-05.com.sun:cluster-iscsi-quorum.btc-ppe-srv10 (not logged in)     iqn.1986-05.com.sun:cluster-iscsi-quorum.btc-ppe-srv9 (not logged in)   testgroup (iSCSI) (ostype: windows):     iqn.1991-05.com.microsoft:btc-ppe-dc.btcppe.netapp.com (logged in on: e0a) </pre>
<p><b>9.</b></p>	<p>Add each node in your cluster to the igroup by using the <code>igroup add</code> command. Make sure you include the cluster node name in this command (for example, clusterA).</p> <p><b>Example:</b> The <code>igroup add</code> command adds the clusterA nodes to the igroup.</p> <pre> btc-ppe-filer7*&gt; igroup add clusterA iqn.1986-05.com.sun:cluster-iscsi- quorum.btc-ppe-wkstn12 btc-ppe-filer7*&gt; igroup add clusterA iqn.1986-05.com.sun:cluster-iscsi- quorum.btc-ppe-wkstn14 </pre> <p>If you are not sure what the cluster node name is, use the <code>scconf</code> command to determine it.</p> <p><b>Example:</b> This example uses the <code>scconf</code> command to get information on the cluster node name and then add the nodes to the igroup.</p> <pre> btc-ppe-wkstn12#scconf -pvv   grep "Cluster node name" Cluster node name:      btc-ppe-wkstn12 Cluster node name:      btc-ppe-wkstn14 </pre>
<p><b>10.</b></p>	<p>Confirm that the nodes were added by entering the <code>igroup show</code> command:</p> <p><b>Example:</b> The <code>igroup show</code> command now includes the nodes you added in Step 10 when it displays a list of the available igroups.</p> <pre> btc-ppe-filer7*&gt; igroup show <b>clusterA (iSCSI) (ostype: solaris):</b>   iqn.1986-05.com.sun:cluster-iscsi-quorum.btc-ppe-wkstn12 (not logged in)   iqn.1986-05.com.sun:cluster-iscsi-quorum.btc-ppe-wkstn14 (not logged in)   clusteragent (iSCSI) (ostype: default):     iqn.1986-05.com.sun:cluster-iscsi-quorum.btc-ppe-srv10 (not logged in)     iqn.1986-05.com.sun:cluster-iscsi-quorum.btc-ppe-srv9 (not logged in)   testgroup (iSCSI) (ostype: windows):     iqn.1991-05.com.microsoft:btc-ppe-dc.btcppe.netapp.com (logged in on: e0a) </pre>
<p><b>11.</b></p>	<p>Map the LUN to the igroup by entering the <code>lun map</code> command.</p> <p><b>Example:</b> The <code>lun map</code> command maps the LUN <code>clusterA.lun</code> to the clusterA igroup.</p> <pre> btc-ppe-filer7*&gt; lun map /vol/vol0/clusterA.lun clusterA lun map: auto-assigned clusterA=0 </pre>
<p><b>12.</b></p>	<p>Confirm that the LUN is mapped to the igroup by entering the <code>lun show -m</code> command.</p> <p><b>Example:</b> The output from the <code>lun show -m</code> command shows that the LUN <code>clusterA.lun</code> is now mapped to clusterA.</p>

	<pre>btc-ppe-filer7*&gt; lun show -m</pre> <table border="1"> <thead> <tr> <th>LUN path</th> <th>Mapped to</th> <th>LUN ID</th> <th>Protocol</th> </tr> </thead> <tbody> <tr> <td>/vol/vol0/clusterA.lun</td> <td>clusterA</td> <td>0</td> <td>iSCSI</td> </tr> <tr> <td>/vol/testflex/testlun0</td> <td>testgroup</td> <td>0</td> <td>iSCSI</td> </tr> <tr> <td>/vol/vol0/clustervol1</td> <td>clusteragent</td> <td>0</td> <td>iSCSI</td> </tr> </tbody> </table> <p><b>Note</b> If the LUN ID is not 0, you must provide the ID to <code>scconf</code> when you add the quorum device.</p>	LUN path	Mapped to	LUN ID	Protocol	/vol/vol0/clusterA.lun	clusterA	0	iSCSI	/vol/testflex/testlun0	testgroup	0	iSCSI	/vol/vol0/clustervol1	clusteragent	0	iSCSI
LUN path	Mapped to	LUN ID	Protocol														
/vol/vol0/clusterA.lun	clusterA	0	iSCSI														
/vol/testflex/testlun0	testgroup	0	iSCSI														
/vol/vol0/clustervol1	clusteragent	0	iSCSI														
<p><b>13.</b></p>	<p>Exit the storage system command line and return to the cluster node.</p> <p><b>Note:</b> You can use Ctrl-D to exit the storage system command line.</p>																
<p><b>14.</b></p>	<p>On a cluster node on the host, add the device as a quorum device by entering the <code>scconf -aq</code> command. You must specify the following information with this command:</p> <ul style="list-style-type: none"> <li>• The name of the storage system. This must be the name you use to identify the device. It must be unique to the cluster.</li> <li>• The LUN ID, if it is not 0 (by default, the LUN ID is set to 0).The command has the following format:</li> <li>•</li> </ul> <pre>scconf -aq name=netapp,type=netapp_nas,filer=&lt;filer_name&gt;[,lun_id=ID]</pre> <p><b>Note</b> For information on administering a quorum device, see the chapter <i>How to Add a Network Appliance NAS Quorum Device</i> in the <i>Sun Cluster System Administration Guide for Solaris OS</i>, which is provided by Sun Microsystems.</p> <p><b>Example:</b> The following command line sets the LUN clusterA on the storage system dev1 as the quorum device. Because the LUN uses 0 as its ID value, you do not need to specify a value for the LUN ID.</p> <pre># scconf -aq name=netapp,type=netapp_nas,filer=btc-ppe-filer7,lun_id=0</pre>																

## 9. Configuring Network Appliance Storage (NFS) Volume for HA File System.

### 9.1. Creating Volumes on Network Appliance Storage

1. Please configure a NetApp storage system running Data ONTAP 7.2 and with NFS and SnapRestore<sup>®</sup> license keys.
2. Create and export volumes for storing shared application files on the storage:
  - a. Create volumes on the storage as listed below.

```
vol create [VolName] [AggrName] [VolSize]
```

Where

*VolName* identifies the name of the volume to be created  
*AggrName* identifies the name of the aggregate that contains the volume  
*VolSize* identifies the size of volume in KB, MB, or GB

The following example creates two volumes named HAcluster and HAoracle.

First create the volumes by issuing the commands at the NetApp storage system:

```
btc-ppe-filer7*> vol create HAcluster aggr1 1024M
```

```
btc-ppe-filer7*> vol create HAoracle aggr1 1024M
```

Then, edit the /etc/exports file on NetApp storage and add the following entries to that file:

```
/vol/HAcluster -sec=sys,rw=btc-ppe-wkstn12:btc-ppe-wkstn14,root=btc-ppe-wkstn12:btc-ppe-wkstn14
```

```
/vol/HAoracle -sec=sys,rw=btc-ppe-wkstn12:btc-ppe-wkstn14,root=btc-ppe-wkstn12:btc-ppe-wkstn14
```

**Note:** /etc/exports file can be edited with “wfile /etc/exports “on NetApp storage after editing the file Ctrl+C to save the entry and exit. You can create the volume using FilerView<sup>®</sup> where entry to /etc/exports will be done automatically.

3. Execute the following command at the storage system console:

```
btc-ppe-filer7*> exportfs -a
```

**Note:** NetApp recommends using flexible volumes in your database environment. NetApp FlexVol<sup>®</sup> technology pools storage resources automatically and enables you to create multiple flexible volumes on a large pool of disks. This flexibility means you can simplify operations, gain maximum spindle utilization and efficiency, and make changes quickly and seamlessly.

## 9.2. Adding Exported Directories from NAS Device to Sun Cluster

### a. From any cluster node, add the directories.

- If you are using Sun Cluster 3.2, use the following command:

```
# clnasdevice add-dir -d /export/dir1, /export/dir2 myfiler
-d /export/dir1, /export/dir2  Enter the directory or directories that you are adding.
myfiler                        Enter the name of the NAS device containing the directories.
```

For more information about the clnasdevice command, see the clnasdevice (1CL) man page.

- If you are using Sun Cluster 3.1, use the following command:

```
# scnasdir -a -h myfiler -d /vol/DB1 -d /vol/DB2
-a          Add the directory or directories to cluster configuration.
-h myfiler  Enter the name of the NAS device whose directories you are adding.
-d          Enter the directory to add. Use this option once for each directory you are adding.
```

This value must match the name of one of the directories exported by the NAS device.

### b. At the prompt, type the HTTP administrator password.

### c. Confirm that the device has been added to the cluster.

- If you are using Sun Cluster 3.2, use the following command:

```
# clnasdevice list
```

For more information about the clnasdevice command, see the clnasdevice (1CL) man page.

- If you are using Sun Cluster 3.1, use the following command:

```
# scnas -p
```

### d. If you do not use the automounter, mount the directories by performing the following steps:

- On each node in the cluster, create a mount point directory for each NAS directory that you added.

```
# mkdir -p /path-to-mountpoint
path-to-mountpoint  Name of the directory on which to mount the directory
```

- On each node in the cluster, add an entry to the /etc/vfstab file for the mount point.

If you are using your NAS device for Oracle® Real Application Clusters database files, set the following mount options:

- forcedirectio
- noac
- proto=tcp

While mounting NAS directories, select the mount options appropriate to your cluster applications. Mount the directories on each node that will access the directories. Sun Cluster places no additional restrictions or requirements on the options that you use.

**Note:** Mount the exported file system on all the nodes and set the mount at boot option to “yes” in /etc/vfstab.

## 10. Configuring Sun cluster with Network Appliance Storage Volume (NFS) for Applications:

This section demonstrates how to mount the NFS file systems onto Sun cluster nodes. There are two volumes /vol/HAcluster and /vol/HAoracle residing on NetApp storage system as an example. Both these volumes are NFS exported at the NetApp storage. The requirement is to mount these volumes on Sun cluster nodes with proper mount options. In this situation NetApp storage acts as NFS server and Sun cluster nodes act as NFS clients.

All cluster nodes mount the NFS file system from NetApp storage; if an active node fails, the other node will detect the failure and will start the application processes on the standby node, thus maintains the high availability. In this scenario Sun cluster HAStoragePlus will not be used for nfs mounted file system.

It is required to install the respective data service agent package for each application used with Sun cluster. Sun cluster will be used as high-availability service for all the dataservice that Sun supports. Sun provides detail documentation for each data service. To configure with sun cluster, this information can be obtained from <http://docs.sun.com>.

Two main categories of cluster-unaware applications can run in the Sun Cluster version 3.2 environments. The applications include:

- **Failover applications**

Failover applications run on only one node of the cluster at a time. The cluster provides high-availability by providing automatic restart on the same or a different node of the cluster.

Failover services are usually paired with an *application IP address*. This is an IP address that always fails over from node to node along with the application. In this way, clients outside the cluster see a *logical host name* with no knowledge on which node a service is running or even knowledge that a service is running in the cluster.

- **Scalable applications**

Scalable applications involve running multiple instances of an application in the same cluster and making it look like a single service by means of a *global interface* that provides a single IP address and load balancing.

While scalable applications are still off-the-shelf, not every application can be made to run as a scalable application in the Sun Cluster 3.2 software environment. Applications that write data without any type of locking mechanism might work as failover applications but do not work as scalable applications.

### 10.1. Preinstallation Requirements:

Sun Cluster supports a number of data services for various applications. Each application has to be installed and configured separately before we create the resource groups and resources for that application.

Failover services are usually paired with an *application IP address* which is used by clients to access the applications. Clients outside the cluster see a logical host name with no knowledge on which node a service is running

1. Install the Sun Cluster application data service software package by running `scinstall` on each node. Use Option 3.

```
# scinstall -I
```

Provide location of the data service package.

2. Create the following:

- a. Mount the nfs exported file system from NetApp storage system to all hosts of the clusters with proper mount options. Mount at boot option should be “yes” in /etc/vfstab. Provide proper permission and ownership access to the directory under mount point.
- b. Create an entry in /etc/inet/hosts for the logical address you will be configuring with the application server.

```
# vi /etc/inet/hosts
. . .
10.73.68.185    apache-lh
10.73.68.187    oracle-lh
```

In the above entry, apache-lh and oracle-lh correspond to the logical hostnames for the applications; with 10.73.68.185 and 10.73.68.187 corresponding to the IP addresses of the logical hosts, respectively.

- c. Create an entry in /etc/inet/hosts for each of the cluster nodes which are members of the cluster and NetApp filer details.
- d. /rhosts – All the IP addresses present in the network with username (i.e.root) need to be entered in this file.

3. Every application related configuration has to be completed as per the Sun data service documents provided by Sun before you create failover or scalable resource groups.

## 10.2. Creating Application Resource Group:

Following are the generic steps for creating resource groups on Network Appliance storage nfs exported volume with Sun cluster. Other resource groups with failover or scalable applications will be created with different external properties as per the data service documents provided by Sun at <http://docs.sun.com>.

1. Registering, Changing, or Removing a Resource Type

You use the following command syntax to register (-a), change (-c), or remove (-r) a resource type:

```
scrgadm -a -t resource_type_name [-h RT_installed_node_list] \
[-f registration_file_path]
```

```
scrgadm -c -t resource_type_name -h RT_installed_node_list
```

```
scrgadm -r -t resource_type_name
```

This is a sample example of adding a resource type for apache.

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

**Note:** The following example is based on the assumption that the APACHE agent from the Sun Cluster 3.2 Data Services CD has been installed by using the scinstall or pkgadd command and Apache related configuration has been done before creating failover or scalable resource group and resources.

2. Adding, Changing, or Removing Resource Groups

```
scrgadm -a -g RG_name [-h nodelist] [-y property [...]]
```

```
scrgadm -c -g RG_name [-h nodelist] -y property [-y property [...]]
```

```
scrgadm -r -g RG_name
```

Example

```
# scrgadm -a -g apache-rg -h btc-ppe-srv9, btc-ppe-srv10
```

### 3. Adding a LogicalHostname or a SharedAddress Resource

```
scrgadm -a -L -g RG_name [-j resource_name] -l hostnamelist [-n netiflist] \  
[-y property [...]]
```

```
scrgadm -a -S -g RG_name -l hostnamelist [-j resource_name] [-n netiflist] \  
[-X auxnodelist] [-y property [...]]
```

Example

```
# scrgadm -a -L -g apache-rg -l apache-lh
```

### 4. SUNW.HASStoragePlus resource on NetApp storage (nfs) exported volume.

**Note:** HASStoragePlus resource will not be supported with nfs mounted file system. Since nfs mounted file system is used as a shared storage for these cluster nodes and Sun cluster does not support nfs exported file system as HASStoragePlus for automatic mount at the time of failover hence nfs exported filesystem will be mounted on all nodes.

### 5. Adding, Changing, or Removing All Resources

```
scrgadm -a -j resource_name -t resource_type_name -g RG_name [-y property [...]] \  
[-x extension_property [...]]
```

```
scrgadm -c -j resource_name [-y property [...]] [-x extension_property [...]]
```

```
scrgadm -r -j resource_name
```

Example

```
#scrgadm -a -j apache-res -t SUNW.apache -y Port_list=80/tcp -x \  
bin_dir=/global/apache/bin -y Network_resources_used=apache-lh -g apache-rg
```

### 6. Make the Application Service online by typing:

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

Example

```
# scswitch -Z -g apache-rg
```

To see current cluster configuration

```
scrgadm -p[v[v]] [-t resource_type_name] [-g resource_group_name] [-j \  
resource_name]
```

**Note:** The above commands will be applicable to Sun cluster 3.1 as well as Sun cluster 3.2.

## 11. References

Sun Cluster overview	<i>Sun Cluster Overview for Solaris OS</i>
Sun Cluster concepts	<i>Sun Cluster Concepts Guide for Solaris OS</i>
Software installation	<i>Sun Cluster Software Installation Guide for Solaris OS</i>
Network-Attached Storage Devices in a Sun Cluster Environment	<i>Sun Cluster 3.1 - 3.2 With Network-Attached Storage Devices Manual for Solaris</i>
Setting up a LUN	<i>Host Cluster Tool for UNIX® Installation Guide</i>
Exporting volumes for the cluster	<i>Data ONTAP Storage Management Guide</i>

## 12. Acknowledgements

The author would like to thank the following individuals for their contribution to the certification process and technical report:

Network Appliance, Inc.

Michelle Nguyen, Uday Shet, Vasu Subbiah

## 13. Disclaimer

Each environment has its own specific set of requirements and no guarantees can be given that the results presented in this report will work as expected on other platforms. This paper should assist in the research and troubleshooting that may be required in a particular case and serve as a checklist of items to be aware of. Please forward any errors, omissions, differences, new discoveries, and comments about this paper to [bijayakumar.meher@netapp.com](mailto:bijayakumar.meher@netapp.com).