Technical Report

CLI Configuration Processes for NetApp and VMware vSphere Storage Arrays Running Data ONTAP and ESX/ESXi Server

Larry Touchette, NetApp

October 2010 | TR-3880

TABLE OF CONTENTS

LIST OF FIGURES

1 INTRODUCTION

TR-3749: NetApp and VMware vSphere Storage Best Practices reviews the best practices for implementing VMware® vSphere™ with NetApp® unified storage arrays running Data ONTAP® and VMware ESX/ESXi Server. The technical report in hand is a companion piece to TR-3749. It provides the information necessary to configure some of the settings applied by the GUI tools described in TR-3749, optional configurations for network and storage layout, and other settings that the customer might prefer to apply manually or by using a CLI.

2 CONFIGURING SSH ON ESX SERVERS AND NETAPP ARRAYS

2.1 CONFIGURING SSH ON ESX SERVERS AND NETAPP ARRAYS

To simplify CLI administration, NetApp FAS systems and VMware ESX Servers can be configured to allow a single host to execute commands remotely on both systems. This management host must have an SSH client installed and configured.

2.2 FAS SYSTEM SSH CONFIGURATION

To configure SSH access on a NetApp FAS system, follow these steps:

1. Connect to the FAS system console (using SSH, telnet, or console connection).
2. Execute the following commands:

   ```
   secureadmin setup ssh
   options ssh.enable on
   options ssh2.enable on
   ```

3. Log in to the Linux® or VMware system that remotely executes commands on the FAS system as root.
4. Add the Triple DES cipher to the list of available SSH ciphers; this is the only cipher recognized by the NetApp FAS system. Edit the `/etc/ssh/ssh_config` file and edit the Ciphers line to read as follows:

   ```
   Ciphers aes128-cbc, aes256-cbc, 3des-cbc
   ```

5. Generate a DSA host key. On a Linux or VMware ESX Server, use the following command:

   ```
   ssh-keygen –t dsa –b 1024
   ```

   When prompted for the passphrase, do not enter one; instead, press Enter.

   The public key is saved to `/root/.ssh/id_dsa.pub`.

6. Mount the FAS root file system as root.
7. Copy only the key information from the public key file to the FAS system's `/etc/sshd/root/.ssh/authorized_keys` file, removing all information except for the key string preceded by the string `ssh-dsa` and a comment line. See the example at the end of these steps.
8. Test the connectivity from the remote host by issuing the `version` command on the FAS system. It should not prompt for a password:

   ```
   ssh <netapp> version
   NetApp Release 7.2: Mon Jul 31 15:51:19 PDT 2006
   ```

Example of the key for the remote host:

```
ssh-dsa AAAAB3NzaC1kc3MAAABhALVbwVyhtAVoaZukcjSTlRb/REO1/ywbQECtAcHijzdzhEJUz
9Qh96HVEwyZDdah+PTxfyitJCerb+1FAnO65v4WMq6jxPVYto6l5Ib5zxfq2I/hhT/6KPziS3LTZj
```

```
KccwAAABUAjkLMwkpiPmg8Unv4fjCsYYhrSL0AAABgF9NsuZxniOOHHr8tmW5RMX+M6VaH/nlJUzV
XbLiI8+pyCXALQ29Y31uV3SzWTd1VOgjJHgv0GBw8N+rvGSB1r60VqgqgGjSB+ZXAO1EecbnjvLnU
tf0TVQ75D9auagjOAAAAYEJPx8wi9/CaS3dfKJR/tYy7Ja+MrlD/RCOgr22XQP1ydexsfYQxenxzE
xPa/sPfjA45YtcUom+3mieFaQuWHZSNFr8sVJoW3LcF5g/z9Wkf5GwvGGtD/yb6bcsjZ4tjlw==
```

2.3 ESX SYSTEM SSH CONFIGURATION

To configure an ESX server to accept remote commands by using SSH, follow these steps:

1. Log in to the ESX console as root.
2. Run the following commands to enable the SSH services:
   ```
   esxcfg-firewall -e sshServer
   esxcfg-firewall -e sshClient
   ```
3. Change to the SSH server configuration directory:
   ```
   cd /etc/ssh
   ```
4. Edit the configuration file:
   ```
   vi sshd_config
   ```
5. Change the following line from
   ```
   PermitRootLogin no
   ```
   to
   ```
   PermitRootLogin yes
   ```
6. Restart the SSH service:
   ```
   service sshd restart
   ```
7. Create the SSH public key:
   ```
   ssh-keygen -t dsa –b 1024
   ```
   This command outputs content similar to the example at the end of these steps. Retain the default locations and do not use a passphrase.
8. Change to the `.ssh` directory:
   ```
   cd /root/.ssh
   ```
9. Run the following commands:
   ```
   cat id_dsa.pub >> authorized_keys
   chmod 600 authorized_keys
   ```
10. Repeat steps 1 through 9 for each ESX server in the cluster.

Example output:

```
Generating public/private dsa key pair.
Enter file in which to save the key (/home/root/.ssh/id_dsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /home/root/.ssh/id_dsa.
Your public key has been saved in /home/root/.ssh/id_dsa.pub.
The key fingerprint is:
7b:ab:75:32:9e:b6:6c:4b:29:dc:2a:2b:8c:2f:4e:37 root@hostname
Your keys are stored in /root/.ssh.
```

3 MULTIMODE ETHERCHANNELS WITH TRADITIONAL ETHERNET SWITCHES

In TR-3749: NetApp and VMware vSphere Storage Best Practices, section 3, “Storage Network Design and Setup,” describes configuring storage networks with traditional Ethernet switches. For reasons of simplicity and routing efficiency, NetApp recommends a configuration based on single-mode EtherChannels; however, not every customer prefers the single-mode configuration. Thus, this section provides a design based on layered multimode EtherChannels. Both NetApp and VMware support this configuration.

3.1 THE LAYERED MULTIMODE DESIGN

The layered multimode design requires each storage controller to have at least four physical network connections, as depicted in Figure 1. The connections are divided into two multimode (active-active) EtherChannels, or VIFs, with IP load balancing enabled. One virtual interface (VIF) is connected to each of the two switches. These two VIFs are then combined into one single mode (active-passive) VIF. NetApp refers to this configuration as a second-level VIF. This option also requires multiple IP addresses on the storage appliance. You can assign multiple IP addresses to the single-mode VIF by using IP address aliases or by using virtual local area network (VLAN) tagging.

3.2 ADVANTAGES OF USING LAYERED MULTIMODE ETHERCHANNEL

Layered multimode EtherChannel provides the following advantages:

- The EtherChannel IP load balancing policy automatically manages storage controller connection load balancing.
- Data I/O to a single IP is aggregated over multiple links.

3.3 DISADVANTAGES OF USING LAYERED MULTIMODE ETHERCHANNEL

Layered multimode EtherChannel contains the following disadvantages:

- Some switch side configuration is required.
- Some storage traffic will cross the uplink between the two switches.

Figure 1 shows the storage side multimode VIFs.

Figure 1) Storage side multimode VIFs.

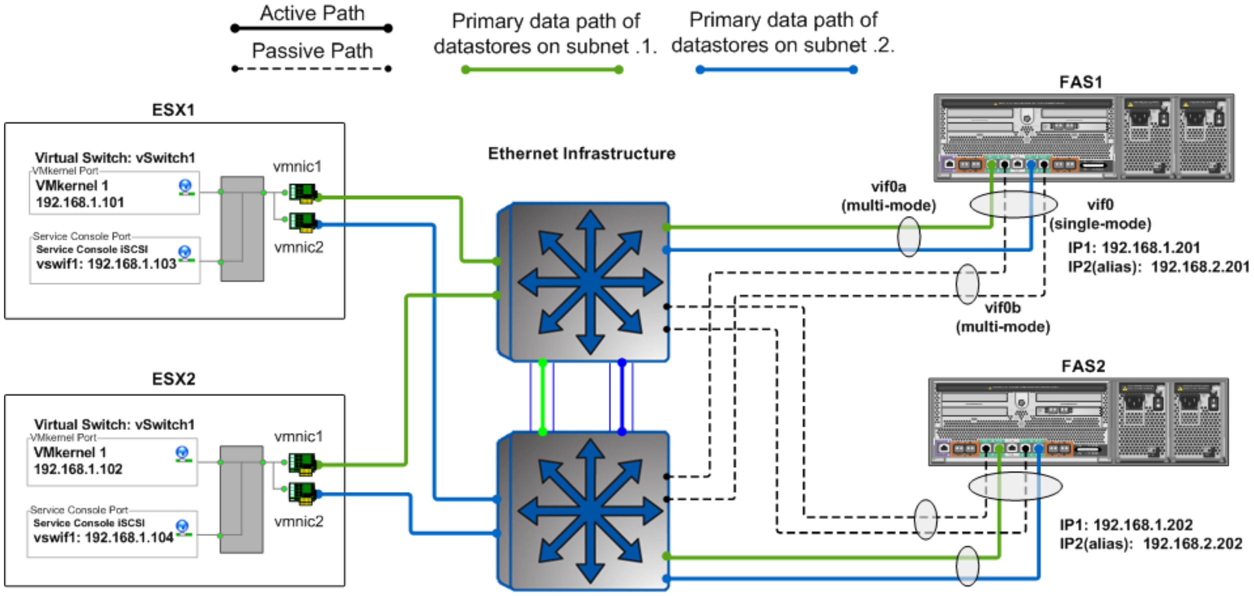

4 CREATING AN ALIGNED VMDK WITH FDISK IN ESX

To create an aligned virtual disk file (virtual machine disk [VMDK]) from the ESX/ESXi console, follow these steps:

1. Log in to the ESX service console.
2. Change to the VM directory:

   ```
   cd /vmfs/volumes/<datastore>/<VM home dir>
   ```

3. View the contents of the VM directory:

   ```
   ls –l
   ```

4. Indentify the number of cylinders in the virtual disk by reading the virtual disk descriptor file. Look for the line `ddb.geometery.cylinders`.

   ```
   cat <Virtual Disk>.vmdk
   ```

5. Run `fdisk` on the virtual disk file (the `–flat.vmdk` file):

   ```
   fdisk ./<Virtual Disk>.vmdk
   ```

6. After accessing `fdisk`, enter extended mode by typing x and pressing Enter.
7. Select the option to set the number of cylinders. Start by typing c and pressing Enter.
8. Enter the number of cylinders that you found in step 4.
9. Type p at the expert command screen to look at the partition table.

   The results should be a table of all zeros.

10. Type r to return to regular mode.
11. To create a new partition, type n and then p when asked for the partition type.
12. Enter 1 for the partition number, 1 for the first cylinder, and press Enter for the last cylinder question to make it use the default value.
13. Type x to access extended mode to set the starting offset.

14. Type b and press Enter, select 1 for the partition and press Enter, then enter 64 and press Enter to set the starting offset.

    Note that the value 64 represents the number of 512-byte blocks used to create a starting offset of 32768kB.

15. Type p to check the partition table.

    If you did this correctly, the top row of the output should display disk geometry, including the starting offset of 64.

16. Type r to return to the regular menu.
17. Type t to set the system type to HPFS/NTF.
18. Enter 7 for the hexcode.
19. Type w to save and write the partition.

    Ignore the warning because this is normal.

20. Start the virtual machine (VM) and run Windows® setup.

    During the installation process, you are prompted that a partition exists. Select this partition in which to format and install Windows.

5 MANUAL STORAGE CONFIGURATIONS FOR FC, FCOE, AND ISCSI

5.1 VERIFY NETAPP HA MODE FOR FC CONFIGURATIONS

NetApp high-availability (HA) arrays ship configured with an option known as cfmode, which controls the behavior of the system’s Fibre Channel (FC) ports if a controller failover occurs. This setting should be set as single system image (SSI). If you are deploying ESX on an older HA array with FC or Fibre Channel over Ethernet (FCoE), then make sure the cfmode is set to SSI.

To verify the current cfmode, follow these steps:

1. Connect to the FAS system console using the SSH, telnet, or console connection.
2. Enter `fcp show cfmode`.

To set the cfmode, follow these steps:

1. Connect to the FAS system console using the SSH, telnet, or console connection.
2. If cfmode needs to be changed, enter FC:

```
set cfmode single_image
```

For more information about the different cfmodes available and the impact of changing a cfmode, see section 8 in the "Data ONTAP Block Management Guide" for the particular version of Data ONTAP that you are running.

5.2 HOST BUS AND CONVERGED NETWORK ADAPTERS

ESX servers and NetApp storage arrays connect to a SAN fabric using host bus adapters (HBAs). Connectivity to FCoE fabrics is enabled through converged network adapters (CNAs). Each HBA/CNA can run as either an initiator (ESX) or a target (NetApp). Each adapter has a unique global address referred to as a World Wide Port Name (WWPN). Each WWPN must be known to configure logical unit number (LUN) access on a NetApp storage array.

Both NetApp and VMware highly recommend that as a best practice, each ESX server should have at least two adapter ports. For more information on VMware FC best practices and recommendations, see VMware Fibre Channel SAN Configuration Guide.

5.3 LUN SIZING FOR VMFS DATASTORES

Virtual Machine File System (VMFS) datastores offer a simple method for provisioning shared storage pools with any storage architecture to implement a design that can address the performance needs of the infrastructure. A common issue is customers overloading very large datastores with too many VMs. In this scenario, the I/O load must be leveled. VMware provides storage VMotion™ as a means to redistribute VM storage to alternative datastores without disruption to the VM. It is common for large VMFS datastores to have reached their I/O performance limit before their capacity limit has been reached.

Although there is no definitive recommendation, a commonly deployed size for a VMFS datastore is somewhere between 300GB and 700GB. The maximum supported LUN size is 2TB. Larger datastores can be created through VMFS spanning. VMFS spanning leverages VMFS extents to concatenate multiple partitions into a single datastore.

Advanced storage technologies such as thin provisioning, which are available with VMware VMDKs and NetApp datastores, can return provisioned but unused storage back to the FAS storage pool for reuse. Unused storage does not include storage that contains data that has been deleted or migrated as part of a storage VMotion process.

5.4 CLUSTER SIZING CONSIDERATIONS WHEN USING LUNS

A VMware cluster is collectively limited to the same maximum number of LUNs as a single ESX server because all ESX/ESXi servers in the cluster must share all storage resources. Currently, the maximum number of LUNs that can be connected to a cluster is 256 LUNs. This limitation typically comes into consideration with VMFS-spanned datastores or raw device mapping (RDM)-based deployments.

Based on LUN limits, the following formula can be used to determine the maximum number of ESX nodes per ESX cluster.

Note: This formula implies that all nodes in a cluster are connected to all shared LUNs.

254/(number of RDMs per VM)/(planned number of VMs per ESX host) = number of ESX nodes in a data center

RDM example:

The formula for 2 RDMs per VM with 20 VMs per ESX server is:

254/2/20 = 6.35 rounded up = 7

In the preceding example, the LUN limit would allow for a maximum of seven ESX hosts in the cluster.

5.5 NETAPP IGROUPS (LUN MASKING)

LUN masking is an authorization process that makes a LUN available to a host or set of hosts in a cluster. On a NetApp array, LUN masking is implemented by assigning HBA addresses to initiator groups (igroups). After an igroup has been defined, LUNs can be assigned to the igroup for access by the host.

Implementation best practices for LUN masking are covered in the storage provisioning section for FC, FCoE, and iSCSI.

5.6 FC, FCOE, AND ISCSI LUN PROVISIONING

When provisioning LUNs for access using FC or iSCSI, the LUNs must be masked so that only the appropriate hosts can connect to them. With a NetApp FAS system, LUN masking is handled by the creation of initiator groups. NetApp recommends creating an igroup for each VMware cluster. NetApp also recommends including the name of the cluster and the protocol type in the name of the igroup (for example, DC1_FC and DC1_iSCSI). This naming convention and method simplify the management of igroups by reducing the total number created. It also means that all ESX servers in the cluster see each LUN at the same ID. Each initiator group includes all of the FC WWPNs or iSCSI-qualified names (IQNs) of the ESX servers in the VMware cluster.

Note: If a cluster uses multiple block-based protocols, separate igroups must be created for each.

For assistance in identifying the WWPN or IQN of the ESX server, select Storage Adapters on the Configuration tab for each ESX server in VMware vCenter™ Server and refer to the SAN Identifier column.

You can create LUNs by using the NetApp LUN wizard in the FAS system console or by using the FilerView® GUI, as shown in the following steps.

1. Log in to FilerView.
2. Select LUNs.
3. Select Wizard.
4. In the Wizard window, click Next.
5. Enter the path.
6. Enter the LUN size.
7. Enter the LUN type (for VMFS select VMware; for RDM select the VM guest OS type).
8. Enter a description and click Next.

The next step in the LUN wizard is LUN masking, which is accomplished by assigning an igroup to a LUN. With the LUN wizard, you can either assign an existing igroup or create a new igroup.

Note: The ESX server expects a LUN ID to be the same on every node in an ESX cluster. Therefore, NetApp recommends creating a single igroup for each cluster rather than a single igroup for each ESX server.

To configure LUN masking on a LUN created in the FilerView GUI, follow these steps:

1. Select Add Group.
2. Select the Use Existing Initiator Group radio button. Click Next and proceed to step 3.

   or

   Select the Create a New Initiator Group radio button. Click Next and proceed to step 4.
3. Select the group from the list and either assign a LUN ID or leave the field blank (the system assigns an ID). Click Next to complete the task.

   Proceed to step 5.
4. Supply the igroup parameters, including name, connectivity type (FC or iSCSI), and OS type (VMware), and then click Next.
5. For the systems that will connect to this LUN, enter the new SAN identifiers or select the known identifiers (WWPN or IQN).
6. Click the Add Initiator button.
7. Click Next to complete the task.

5.7 CONNECTING FC AND FCOE DATASTORES

The Fibre Channel service is the only storage protocol that is running by default on the ESX server. NetApp recommends that each ESX server has two FC HBA ports available for storage path redundancy.

To connect to FC LUNs provisioned on a FAS system, follow these steps:

1. Open vCenter Server.
2. Select an ESX host.

3. In the right pane, select the Configuration tab.
4. In the Hardware box, select the Storage Adapters link.
5. In the upper-right corner, select the Rescan link.
6. Repeat steps 1 through 5 for each ESX server in the cluster.

Selecting Rescan forces the rescanning of all HBAs (FC, FCoE, and iSCSI) to discover changes in the storage available to the ESX server.

To add a LUN as a datastore, follow these steps:

1. Open vCenter Server.
2. Select an ESX host.
3. In the right pane, select the Configuration tab.
4. In the Hardware box, select the Storage link and then click Add Storage to open the Add Storage wizard.
5. Select the Disk/LUN radio button and click Next.
6. Select the LUN to use and click Next.
7. Enter a name for the datastore and click Next.
8. Select the block size, click Next, and click Finish.

The default block size of a virtual machine file system is 1MB. This block size supports storing virtual disk files up to a maximum of 256GB in size. If you plan to store virtual disks larger than 256GB in the datastore, you must increase the block size to be greater than the default.

5.8 ENABLE ISCSI COMMUNICATIONS

To enable iSCSI communications, follow these steps:

1. Open vCenter Server.
2. Select an ESX host.
3. In the right pane, select the Configuration tab.
4. In the Configuration tab, left pane, select Security Profile.
5. In the right pane, select the Properties link to open the Firewall Properties window.
6. Select the Software iSCSI Client checkbox and then click OK to close the Firewall Properties window.

5.9 CONNECTING TO ISCSI TARGETS

To connect to iSCSI targets, follow these steps:

1. Open vCenter Server.
2. Select an ESX host.
3. In the right pane, select the Configuration tab.
4. In the right pane, Hardware box, select Storage Adapters.
5. Highlight the iSCSI adapter and click the Properties link in the Details box.
6. Select the Dynamic Discovery tab in the iSCSI Initiator Properties box.
7. Click Add and enter the IP address of the iSCSI-enabled interface on the NetApp FAS system.
8. For an additional layer of security, select the CHAP tab to configure CHAP authentication. NetApp recommends setting up and verifying iSCSI access before enabling CHAP authentication.

5.10 RESTRICTING ISCSI TARGETS TO PREFERRED INTERFACES

By default, NetApp storage arrays provide iSCSI access over every network interface. This default configuration might not be optimal because it can lead to conditions where ESX servers attempt to communicate to interfaces that are unreachable. NetApp recommends that you disable iSCSI on NetApp network interfaces over which you do not want to send iSCSI traffic.

Data ONTAP allows this filtering to be accomplished either on a host-by-host basis using iSCSI access lists or on a global basis by unbinding iSCSI to a specific interface or set of interfaces. NetApp recommends that you use one of these two methods to configure iSCSI access restrictions.

Host restricted iSCSI access lists currently require each IQN of an ESX server to be configured on the array. This process is more granular and might lead to additional tasks each time a new host is introduced into the data center.

To configure iSCSI access lists, follow these steps:

1. Connect to the FAS system console (using SSH, telnet, or console connection).
2. To create an iSCSI access list, type:
   ```
   iscsi interface accesslist add <ESX iqn address>
   ```
3. Repeat step 2 for each ESX host in the data center.
4. To verify the iSCSI access list, type:
   ```
   iscsi interface accesslist show
   ```

Globally disabling iSCSI traffic on a set of network interfaces is less granular than iSCSI access lists; however, it is much simpler to configure.

To globally disable iSCSI traffic on a set of network interfaces, follow these steps:

1. Connect to the FAS system console (using SSH, telnet, or console connection).
2. To disable iSCSI on an interface, type:
   ```
   iscsi interface disable <interface hw address>
   ```
3. To verify the iSCSI bindings, type:
   ```
   iscsi interface show
   ```

6 VMWARE NATIVE MULTIPATHING

VMware ESX Servers ship with a native multipathing solution for FC, FCoE, and iSCSI storage networks, which enables high-performance data access and enhanced storage resiliency. With the release of ESX and ESXI 4.0, VMware has introduced the concept of a pluggable storage architecture (PSA), which introduced several new concepts to its native multipathing policy (NMP). This section reviews leveraging the storage array type plug-in (SATP) and the path selection plug-in (PSP) along with the asymmetric logical unit access (ALUA) protocol.

6.1 DEFAULT NMP SETTINGS

Connecting a NetApp array to an ESX 4.x server results in the array being identified as an active-active storage controller and the VMware native multipathing path selection policy applying the fixed multipathing policy. This configuration is identical to the default behavior with ESX 3.5.

Deployments that leverage the fixed multipathing policy are required to manually identify and set the I/O to traverse the primary FC paths. In addition, users of this configuration are required to manually load balance I/O across the primary paths. The NetApp ESX host utilities (EHU) can automate this process for environments that prefer the NMP fixed PSP. Note that the EHU has been superseded by the Virtual Storage Console (VSC) described in TR-3749: NetApp and VMware vSphere Storage Best Practices.

For deployments that prefer a complete plug-and-play architecture, enable ALUA on the NetApp storage array and configure the round robin PSP.

6.2 ENABLING ALUA

NetApp and VMware support the ALUA protocol. ALUA allows for the autonegotiation of paths between SCSI target devices and target ports, enabling dynamic reconfiguration. Enabling ALUA on NetApp initiator groups results in a more dynamic, or plug-and-play-like, architecture.

Note: ALUA is supported with ESX as well as ESXI for FC and FCoE. Support for iSCSI is not required because iSCSI addresses this functionality natively within the protocol.

ALUA is enabled on ESX 4.0 by default.

To enable ALUA on a NetApp storage array, follow these steps:

1. Log in to the NetApp console.
2. From the storage appliance console, run:

   ```
   igroup set <igroup-name> alua yes
   ```
3. Repeat step 2 for each LUN accessed by ESX.
4. Results can be verified by running:

   ```
   igroup show -v <igroup-name>
   ```

6.3 DEFAULT NMP SETTINGS WITH ALUA ENABLED

Connecting a NetApp array to an ESX 4.0 server with ALUA enabled results in the array and server being able to negotiate which paths are primary for I/O and which paths should be used for failover. When ALUA is enabled, the array is identified as an ALUA-enabled storage controller, and the VMware native multipathing path selection policy applies the most recently used (MRU) multipathing policy.

Deployments that leverage ALUA along with the MRU multipathing policy are required to manually load balance I/O across the primary paths. The result of only enabling ALUA is a reduction in some of the configuration requirements. For deployments that prefer a complete plug-and-play architecture, enable ALUA on the NetApp storage array and configure the round robin PSP.

The VMware round robin path selection policy enables link bandwidth aggregation because it sends I/O packets down each of the preferred active paths. This PSP is optimal when combined with ALUA and LUNs that do not implement a per-LUN queue depth limit. LUNs served by Data ONTAP do not implement such limits.

When AULA is enabled on LUNs that are accessed by both ESX/ESXi3 and ESX/ESXi4 hosts, NetApp recommends using the fixed PSP. VMware KB1010713 provides additional information. At the time of this publication, VMs running MSCS are not supported with the RR PSP.

6.4 MANUALLY CONFIGURING THE ROUND ROBIN PSP

There are two ways to configure a PSP. NetApp recommends setting the ESX system default PSP for the VMware default ALUA SATP to use the round robin PSP. Alternatively, you can manually manage datastore and LUN policies inside the virtual infrastructure client, as was done in ESX 3.5.

6.5 SETTING THE DEFAULT PSP FOR ALUA TO ROUND ROBIN

To set the default PSP for ALUA to round robin, follow these steps:

1. Connect to the CLI of an ESX or ESXi server.
2. From the console, run:

   ```
   esxcli nmp satp setdefaultpsp -–psp <PSP type> -–satp <SATP type>
   ```

   Available PSP types in vSphere are:

   ```
   VMW_PSP_RR
   VMW_PSP_FIXED
   VMW_PSP_MRU
   Available SATP types for NetApp arrays are:
   VMW_SATP_DEFAULT_AA
   VMW_SATP_ALUA
   ```

   Example of executing this command:

   ```
   esxcli nmp satp setdefaultpsp --psp VMW_PSP_RR --satp VMW_SATP_ALUA
   ```

3. To verify the results of this command, type

   ```
   esxcli nmp satp list:
   ```

4. Repeat steps 1 through 3 for each ESX or ESXi server.

6.6 MANUALLY SETTING THE PSP FOR A DATASTORE

1. Open vCenter Server.
2. Select an ESX server.
3. In the right pane, select the Configuration tab.
4. In the Hardware box, select Storage.
5. In the Storage box, highlight the storage and select the Properties link.
6. In the Properties dialog box, click the Manage Paths button.
7. Set the multipathing policy to round robin.

6.7 MANUALLY SETTING THE PSP FOR A LUN

An alternative method for setting the preferred path for multiple LUNs is available in vCenter Server.

To set the path, follow these steps:

1. Open vCenter Server.
2. Select an ESX server.
3. In the right pane, select the Configuration tab.
4. In the Hardware box, select Storage Adapters.
5. In the Storage Adapters pane, select a host bus adapter.
6. Highlight all of the LUNs to configure.

7. Right-click the highlighted LUNs and select Manage Paths.
8. In the Manage Paths window, set the multipathing policy and preferred path for all of the highlighted LUNs.

7 MANUAL STORAGE CONFIGURATIONS FOR NFS

7.1 INCREASING THE NUMBER OF NFS DATASTORES

By default, VMware ESX is configured to support up to eight Network File System (NFS) datastores; however, this limit can be increased to 64 in order to meet the needs as the virtual infrastructure grows. While the maximum number of NFS datastores (64) is less than what is available with VMFS datastores (256), this difference is offset by the density available to NetApp NFS datastores.

To make sure of availability, NetApp recommends that you increase the maximum number of datastores available when deploying an ESX server. Preconfiguring this setting makes sure that NFS datastores can be added dynamically at any time without disruption or effort.

To make this change, follow these steps from within the virtual infrastructure client:

1. Open vCenter Server.
2. Select an ESX host.
3. In the right pane, select the Configuration tab.
4. In the Software box, select Advanced Configuration.
5. In the pop-up window, left pane, select NFS.
6. Change the value of NFS.MaxVolumes to 64.
7. In the pop-up window, left pane, select Net.
8. Change the value of Net.TcpIpHeapSize to 30.
9. Change the value of Net.TcpIpHeapMax to 120.
10. Repeat steps 1 through 9 for each ESX server.

7.2 FILE SYSTEM SECURITY

NetApp storage arrays allow customers to set the security style of each flexible volume (or file system) to use UNIX® permissions or NT file system (NTFS) permissions. File system security can be mixed and matched with share or export security. As an example, a UNIX share (or export) can allow access to a file system with NTFS permissions and vice versa. In addition, security style can also be made on a file-by-file basis using the MIXED permissions setting.

For VMware deployments, NetApp highly recommends setting the security style of all datastores to UNIX. The security setting of the root volume becomes the security setting when you create a new volume.

Customers running VMware on NFS commonly want to access their datastores from Windows systems in order to complete administrative functions. With this use case in mind, set the volume security style to UNIX and make sure that the FAS user mapping is set up correctly for Windows users to access this data. For more information on this subject, review the section “File Sharing Between NFS and CIFS” in the "Data ONTAP File Access and Protocol Management Guide" for the particular version of Data ONTAP that you are running.

If you need to change the file system security type, follow these steps:

1. Log in to the NetApp console.
2. From the storage appliance console, run:

```
vol options <vol-name> no_atime_update on
```

3. From the storage appliance console, run:

```
qtree security <volume path> UNIX
```

4. Repeat steps 2 and 3 for each NFS-accessed volume.

7.3 ESX NFS TIMEOUT SETTINGS

When connecting to NFS datastores, NetApp recommends adjusting selective NFS options related to connection monitoring and resiliency.

For optimal availability with NFS datastores, NetApp recommends making the following changes on each ESX 4.0 host:

1. Open vCenter Server.
2. Select an ESX host.
3. In the right pane, select the Configuration tab.
4. In the Software box, select Advanced Configuration.
5. In the pop-up window, left pane, select NFS.
6. Change the value of NFS.HeartbeatFrequency to 12.
7. Change the value of NFS.HeartbeatMaxFailures to 10.
8. Repeat steps 1 through 7 for each ESX server.

7.4 NFS TCP WINDOW SIZE

In previous versions of the NetApp best practice guide, NetApp recommended increasing the Transmission Control Protocol (TCP) window size (or Data ONTAP option `nfs.tcp.recvwindowsize`). With Data ONTAP 7.2.4 and later releases, there is no need to modify this setting.

7.5 CONNECTING NFS DATASTORES

To create a file system for use as an NFS datastore, follow these steps:

1. Open FilerView (for example, http://<filer>/na_admin).
2. Select Volumes.
3. Select Add to open the Volume wizard. Complete the wizard.
4. From the FilerView menu, select NFS.
5. Select Add Export to open the NFS Export wizard. Complete the wizard for the newly created file system, granting read/write and root access to the VMkernel address of all ESX hosts that will connect to the exported file system.
6. Open vCenter Server.
7. Select an ESX host.
8. In the right pane, select the Configuration tab.
9. In the Hardware box, select the Storage link.
10. In the upper-right corner, click Add Storage to open the Add Storage wizard.
11. Select the Network File System radio button and click Next.
12. Enter a name for the storage appliance, export, and datastore, then click Next.
13. Click Finish.

8 MANUALLY SETTING VOLUME AUTOGROW

Volume autosize is a policy-based space management feature of Data ONTAP that allows a volume to grow in defined increments up to a predefined limit when the volume is nearly full. For VMware environments, NetApp recommends setting this value to on. Doing so requires setting the maximum volume and increment size options.

To enable these options, follow these steps:

1. Log in to NetApp console.
2. Set volume autosize policy:

   ```
   vol autosize <vol-name> [-m <size>[k|m|g|t]] [-i <size>[k|m|g|t]] on
   ```

   For example, to set a maximum volume size of 500 gigabytes with an increment grow size of 10 gigabytes on a volume called `vsphere1`, use the following command:

   ```
   vol autosize vsphere1 –m 500g –i 10g on
   ```

Snapshot™ autodelete is a policy-based space-management feature that automatically deletes the oldest Snapshot copies on a volume when that volume is nearly full. For VMware environments, NetApp recommends setting this value to delete Snapshot copies at 5% of available space. In addition, you should set the volume option to have the system attempt to grow the volume before deleting Snapshot copies.

To enable these options, follow these steps:

1. Log in to NetApp console.
2. Set Snapshot autodelete policy:

   ```
   snap autodelete <vol-name> commitment try trigger volume target_free_space
   5 delete_order oldest_first
   ```

3. Set volume autodelete policy:

   ```
   vol options <vol-name> try_first volume_grow
   ```

LUN fractional reserve is a policy that is required when you use NetApp Snapshot copies on volumes that contain VMware LUNs. This policy defines the amount of additional space reserved to guarantee LUN writes if a volume becomes 100% full. For VMware environments where volume autosize and Snapshot autodelete are in use and you have separated the swap, pagefile, and other transient data onto other LUNs and volumes, NetApp recommends setting this value to 0%. Otherwise, leave this setting at its default of 100%.

To enable this option, follow these steps:

1. Log in to NetApp console.
2. Set volume Snapshot fractional reserve:

   ```
   vol options <vol-name> fractional_reserve 0
   ```

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